TECHNICAL REPORT

TR-350

Ethernet Services using BGP MPLS Based Ethernet VPNs (EVPN)

Issue: 2
Issue Date: TBD 2017

Notice

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	9 Nov. 2015	11 Nov. 2015	Rao Cherukuri, Juniper Sriganesh Kini, Ericsson	Original
bbf2016.691.00	TBD	TBD	Rao Cherukuri, BBF Distinguished Fellow Steven Blake, Ericsson	Added Sec. 13 EPL and EVPL
bbf2016.691.01	TBD	TBD	Rao Cherukuri, BBF Distinguished Fellow Steven Blake, Ericsson	Correct [51] URL, R-81; fix editorial nits.
bbf2016.691.02	TBD	TBD	Rao Cherukuri, BBF Distinguished Fellow Steven Blake, Ericsson	Same as .01 with all changes accepted.
WT-350 ELine Modifications from bbf2016.691.02-Cherukuri-DF-bbf2016.691.doc	TBD	TBD	Rao Cherukuri, BBF Distinguished Fellow Steven Blake, Ericsson	Same as .02 with changes for Bandwidth Policy and updated figures for Eline.
WT-350 ELine final Modifications from bbf2016.691.02-Cherukuri-DF-bbf2016.691.docx	TBD	TBD	Rao Cherukuri, BBF Distinguished Fellow	Same as above with changes for ETree.
WT-350 ELine final Modifications from bbf2016.691.02-Cherukuri-DF-bbf2016.691.docx	TBD	TBD	Rao Cherukuri, BBF Distinguished Fellow	Proposed SB Candidate
WT-350i2ELineETr	TBD	TBD	Rao Cherukuri, BBF Distinguished Fellow	SB Text

ee.SBCandidate-Sinicrope-RnTWAD				

Comments or questions about this Broadband Forum Technical Report should be directed to help@broadband-forum.org.

Editor	Rao Cherukuri	BBF Distinguished Fellow
	Steven Blake	Ericsson
Routing and Transport WA Director	David Sinicrope	Ericsson

TABLE OF CONTENTS

List of Figures

Executive Summary

Carrier Ethernet provides extensions to Ethernet, enabling telecommunications network providers to provide Ethernet services to customers and to utilize Ethernet technology in their networks.

Carrier Ethernet services are being used in Broadband access networks, enterprise networks and backhaul networks. Providing Carrier Ethernet services using MPLS network infrastructure is generating revenue opportunities for global carriers, driven by customer demand for higher bandwidth connectivity. Though TR-224 describes the architecture for solutions to implement Carrier Ethernet services using an MPLS network, the VPLS based solution has a number of limitations when it comes to redundancy, multicast optimization and provisioning simplicity. It does not address requirements such as multi-homing with all-active forwarding, load balancing, policy-based control and control plane based MAC learning. Service interface requirements for data-center interconnects are also not addressed by TR-224.

This document provides technical architecture and equipment requirements to implement the Carrier Ethernet services using BGP MPLS-based EVPNs in order to overcome the limitations of VPLS and address the additional requirements. By specifying a common technical architecture, common equipment requirements and common set of feature options, this document promotes multi-vendor interoperability.

1 Purpose and Scope

1.1 Purpose

Carrier Ethernet provides extensions to Ethernet enabling telecommunications network providers to provide Ethernet services to customers and to utilize Ethernet technology in their networks. Service providers are deploying Carrier Ethernet services around the globe, in large part, because Carrier Ethernet has compelling capabilities such as standardized service definitions as well as improved scalability, reliability, QoS, and manageability.

Carrier Ethernet services are being used in Broadband access networks, enterprise networks and backhaul networks. The integration of Ethernet into MPLS network infrastructure is generating revenue opportunities for global carriers, driven by customer demand for higher bandwidth connectivity. This document provides a technical architecture and equipment requirements implementing the specified Ethernet services using BGP MPLS-based Ethernet VPNs (EVPNs) in IP/MPLS network.

New Ethernet service applications require capabilities such as: multi-homing with all-active forwarding; load balancing; policy based control, and control plane MAC learning. TR-224 [4] and TR-178 [2] based solutions do not provide these features; solutions based on BGP MPLS EVPNs do.

By specifying a common technical architecture, common equipment requirements and common set of feature options, this document promotes multi-vendor interoperability. This document may be used as a basis for conformance testing.

1.2 Scope

This document defines a reference architecture for Carrier Ethernet services using BGP MPLS-based Ethernet VPN mechanisms:

- Ethernet multipoint to multipoint (E-LAN)
- Ethernet point to point (E-Line)
- Ethernet point to multipoint (E-Tree) including E-Tree*
- Ethernet access to support wholesale access service
- Control, OAM, QoS, reliability and scalability for the MPLS network
- Support Ethernet service capabilities specified in RFC 7209 [37]

This document specifies how to implement the Ethernet services layer. It does not specify the service layer itself. Ethernet Control and OAM protocols will be transparently transported, except for cases where Layer 2 control protocol processing is required per-service definition.

This document includes the Carrier Ethernet E-LAN, E-Line, and E-Tree (including E-Tree*) services types using BGP MPLS EVPNs.

In order to support Carrier Ethernet services across multiple networks, the scope of this document includes the following:

- Attachment circuits (AC) providing a user-to-network interface complying with the MEF UNI are supported.
- Ethernet attachment circuits for multi-service broadband access and aggregation (i.e., TR-101/TR-178) are supported.
- Additional Ethernet service capabilities of BGP MPLS-based EVPNs (e.g. multi-homing with all-active forwarding, load balancing, policy based control, control based MAC learning, etc) are supported.
- Support for interworking with TR-224.
- To support Carrier Ethernet across multiple SP networks, the specification addresses multi autonomous systems which preserves end-to-end capabilities (e.g., OAM, QoS and protection etc).
- Cases where the UNI-N functions are or are not collocated with the PE are supported.

TR-350 provides technical architecture and equipment requirements implementing MEF Carrier Ethernet services with BGP MPLS EVPNs. EVPNs architecture and protocols are based on BGP/MPLS IP VPNs , which supports multiple domains. This capability is used to support connectivity between service endpoints (e.g. MEF UNIs) connected to different networks or operators.

TR-350 does not use architecture or connectivity models of Carrier Ethernet using MEF 26.1 [46].

2 References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found be in RFC 2119 [8].

MUST	This word, or the term "REQUIRED", means that the definition is an absolute requirement of the specification.
MUST NOT	This phrase means that the definition is an absolute prohibition of the specification.
SHOULD	This word, or the adjective "RECOMMENDED", means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course.
SHOULD NOT	This phrase, or the phrase "NOT RECOMMENDED" means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label.
MAY	This word, or the adjective "OPTIONAL", means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

2.2 References

The following references are of relevance to this Technical Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Technical Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.
A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

Document	Title	Source	Year
[1] TR-101	*Migration to Ethernet-Based Broadband Aggregation*	BBF	2011
[2] TR-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[3] TR-221	*Technical Specifications for MPLS in Mobile Backhaul Networks*	BBF	2011

[4] TR-224	*Technical Specification for MPLS in Carrier Ethernet Networks*	BBF	2014
[5] IEEE 802.3	*IEEE Standard Ethernet*	IEEE	2012
[6] IEEE 802.1Q	*IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks*	IEEE	2011
[7] RFC 1195	*Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*	IETF	1990
[8] RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[9] RFC 2328	*OSPF Version 2*	IETF	1998
[10] RFC 3209	*RSVP-TE: Extensions to RSVP for LSP Tunnels*	IETF	2001
[11] RFC 3270	*Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*	IETF	2002
[12] RFC 3473	*Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*	IETF	2003
[13] RFC 3478	*Graceful Restart Mechanism for Label Distribution Protocol*	IETF	2003
[14] RFC 3564	*Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*	IETF	2003
[15] RFC 3623	*Graceful OSPF Restart*	IETF	2003
[16] RFC 3630	*Traffic Engineering (TE) Extensions to OSPF Version 2*	IETF	2003
[17] RFC 5306	*Restart Signaling for Intermediate System to Intermediate System (IS-IS)*	IETF	2004
[18] RFC 4090	*Fast Reroute Extensions to RSVP-TE for LSP Tunnels*	IETF	2005
[19] RFC 4124	*Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*	IETF	2005
[20] RFC 4206	*Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*	IETF	2005
[21] RFC 4364	*BGP/MPLS IP Virtual Private Networks (VPNs)*	IETF	2006

[22]	RFC 4379	*Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*	IETF	2006
[23]	RFC 4761	*Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*	IETF	2007
[24]	RFC 4762	*Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*	IETF	2007
[25]	RFC 5036	*LDP Specification*	IETF	2007
[26]	RFC 5150	*Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)*	IETF	2008
[27]	RFC 5151	*Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*	IETF	2008
[28]	RFC 5283	*LDP Extension for Inter-Area Label Switched Paths (LSPs)*	IETF	2008
[29]	RFC 5286	*Basic Specification for IP Fast Reroute: Loop-Free Alternates*	IETF	2008
[30]	RFC 5305	*IS-IS Extensions for Traffic Engineering*	IETF	2008
[31]	RFC 5586	*MPLS Generic Associated Channel*	IEIF	2009
[32]	RFC 5880	*Bidirectional Forwarding Detection (BFD)*	IETF	2010
[33]	RFC 5881	*Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)*	IETF	2010
[34]	RFC 5884	*Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*	IETF	2010
[35]	RFC 6424	*Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*	IETF	2011
[36]	RFC 6790	*The Use of Entropy Labels in MPLS Forwarding*	IETF	2012
[37]	RFC 7209	*Requirements for Ethernet VPN (EVPN)*	IETF	2014
[38]	RFC 7387	*Framework for Ethernet Tree (E-Tree) Service over a Multiprotocol Label Switching (MPLS) Network*	IETF	2014
[39]	RFC 7432	*BGP MPLS Based Ethernet VPN*	IETF	2015
[40]	G.8013/Y.1731	*OAM functions and mechanisms for Ethernet based networks*	ITU-T	2013
[41]	MEF 6.1	*Ethernet Services Definitions - Phase 2*	MEF	2008

[42]	MEF 10.2	*Ethernet Services Attributes - Phase 2*	MEF	2009
[43]	MEF 30	*Service OAM Fault Management Implementation Agreement*	MEF	2011
[44]	MEF 22.1	*Mobile Backhaul Phase 2 Implementation Agreement*	MEF	2012
[45]	MEF 23.1	*Carrier Ethernet Class of Service – Phase 2*	MEF	2012
[46]	MEF 26.1	*External Network Network Interface (ENNI) – Phase 2*	MEF	2012
[47]	MEF 35	*Service OAM Performance Monitoring Implementation Agreement*	MEF	2012
[48]	MEF 6.1.1	*Layer 2 Control Protocol Handling Amendment to MEF 6.1*	MEF	2012
[49]	MEF 10.3	*Ethernet Services Attributes - Phase 3*	MEF	2013
[50]	MEF 6.2	*EVC Ethernet Services Definitions - Phase 3*	MEF	2014
[51]	MEF 45	*Multi-CEN L2CP*	MEF	2014
[52]	draft-ietf-bess-evpn-etree	*E-TREE Support in EVPN & PBB-EVPN*	IETF	2017
[53]	draft-ietf-bess-evpn-vpws	*Virtual Private Wire Support in EVPN*	IETF	2017

2.3 Definitions

The following terminology is used throughout this Technical Report.

AGN	An aggregation node (AGN) is a node which aggregates several access nodes (ANs).
AN	An access node is a node which processes customer frames or packets at Layer 2 or above. This includes but is not limited to DSLAMs or OLTs (in case of (G)PON deployments).
E-Line	A service connecting two customer Ethernet ports over a WAN.
E-LAN	A multipoint service connecting a set of customer endpoints, giving the appearance to the customer of a bridged Ethernet network connecting the sites.
E-Tree	A rooted multipoint Ethernet Virtual Connection. The E-Tree service type can support one or multiple Root UNIs (see Section 9.3/MEF 6.2 [50])

E-Tree*	Partially implementing MEF multipoint service connecting only one root and a set of leaves, but preventing inter-leaf communication. See details in TR-221. Note: Ethernet Tree (E-Tree) service type is specified in Section 6.3/MEF 6.1 [41]. The Appendix in TR-221 modifies the E-Tree service type which is used in different services. The modified E-Tree* service type is used in both Ethernet Private Tree service and Ethernet Virtual Private Tree Service specified in Section 14 .

2.4 Abbreviations

This Technical Report uses the following abbreviations:

AC	Attachment Circuit
AGN	Aggregation Node
AN	Access Node
BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
BNG	Broadband Network Gateway
CE	Customer Edge
CoS	Class of Service
CV	Connectivity Verification
EPL	Ethernet Private Line
EP-LAN	Ethernet Private-LAN
EVC	Ethernet Virtual Connection
EVPL	Ethernet Virtual Private Line
EVP-LAN	Ethernet Virtual Private – LAN
FD	Frame Delay
FRR	Fast Re-route
FLR	Frame Loss Ratio
H-VPLS	Hierarchal Virtual Private LAN Service
IETF	Internet Engineering Task Force
IP	Internet Protocol
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
L2VPN	Layer 2 Virtual Private Network
LAN	Local Area Network
LER	Label Edge Router
LFA	Loop Free Alternate
LSP	Label Switched Path

LSR	Label Switch Router
MAC	Medium Access Control
MEF	Metro Ethernet Forum
MPLS	Multi Protocol Label Switching
OAM	Operations, Administration and Management
OAMPDU	OAM Protocol Data Unit
P	Provider
PE	Provider Edge
PSN	Packet Switched Network
PW	Pseudowire
QoS	Quality of Service
RFC	Request for Comments
RSVP-TE	Resource ReSerVation Protocol with Traffic Engineering extensions
SLA	Service Level Agreement
TE	Traffic Engineering
T-LDP	Targeted Label Distribution Protocol
TLV	Type/Length/Value
TR	Technical Report
UNI	User to Network Interface
UDP	User Datagram Protocol
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
VPWS	Virtual Private Wire Service
WG	Working Group

3 Technical Report Impact

3.1 Energy Efficiency

TR-350 has no impact on energy efficiency.

3.2 IPv6

Carrier Ethernet services operate at Layer 2 and therefore the network is agnostic to IPv6 user traffic. The IPv6 header DSCP field is assumed to be mapped to the Ethernet P bits by the service user.

IPv6 addressing may appear in its respective places in control, OAM, and management protocols, e.g., node ids, FECs, and loopback addresses, etc.

TR-350 has no direct impact on IPv6.

3.3 Security

Security requirements are specified for each service in respective sections.

3.4 Privacy

Any issues regarding privacy are not affected by TR-350.

4 Carrier Ethernet services

Ethernet is now being used as both a transport technology and a service delivery architecture. The MEF Carrier Ethernet specifies Ethernet service type, service attributes, QoS and SLA. The service type includes point-to-point (E-Line), point-to-multipoint (E-Tree), multipoint-to-multipoint (E-LAN), and E-Access. The service definition includes both port-based and VLAN-based service identification.

TR-224 refers to MEF 6.1 and MEF 10.2. TR-350 uses the backward compatible subset of the revised specifications MEF 6.2 [50] and MEF 10.3 [49] to achieve the equivalent function. This makes WT-350 compatible with TR-224.

The MEF also defined Carrier Ethernet as a ubiquitous, standardized, carrier-class service and network defined by attributes that distinguish Carrier Ethernet from familiar LAN-based Ethernet.

4.1 Carrier Ethernet Requirements

Service providers worldwide are migrating their existing networks to deliver Carrier Ethernet services to enterprises, businesses & residential end-users. The attributes are as follows:

1. Standardized Services
   - Support E-Line, E-LAN and E-Tree service types as defined by MEF
   - No changes to customer LAN equipment or networks and accommodates existing network connectivity such as, time-sensitive, TDM traffic and signaling
   - Wide choice and granularity of bandwidth and quality of service options
2. Security
3. Scalability
   - The ability to support millions of Ethernet Virtual Connection (EVC) services for enterprise and residential users
   - Scalability of bandwidth from 1 Mbps to 10 Gbps and beyond, in granular increments
4. Reliability
   - The ability for the network to detect & recover from faults quickly
   - Fast network convergence
5. Quality of Service
   - Service Level Agreements (SLAs) that deliver end-to-end performance
   - Traffic profile enforcement per-EVC
   - Hierarchical queuing
6. Service Management
   - Minimize network touch points in provisioning
   - Standards-based OAM to support SLA

5 Layer 2 Ethernet VPNs in MPLS Networks

MPLS has for a long time been defined as a convergence technology, one that will allow service providers to bring together their disparate networks and leverage features like traffic engineering, hierarchal QoS, and service interworking.

Provider Provisioned Virtual Private Networks (PPVPN) now dominates the IP-VPN services market and is projected for significant growth. Many service providers have provided Ethernet VPN services using Virtual Private LAN Services (VPLS) as an alternative that allows enterprises to manage their own routing.

TR-224 uses VPLS to support Ethernet LAN services in IP/MPLS networks. A VPLS PE emulates an Ethernet bridge (IEEE 802.1Q [6]) and performs MAC learning in the data plane. New applications using Ethernet services require capabilities such as: multi-homing with all-active forwarding, load balancing, policy based control, and control plane MAC learning. To support these capabilities IETF developed BGP MPLS-based Ethernet VPNs (EVPNs). TR-224 based solutions do not provide these features.

When an Ethernet multipoint service is provided using EVPN, control-plane-based remote MAC learning is used over the MPLS core (PE to PE) network. MAC learning between PE and CE is done in the data plane. EVPN is designed to handle multi-homing, and per-flow load balancing. The EVPN technology uses MP-BGP over an MPLS network. The technology is similar to BGP MPLS-based IP VPNs (RFC 4364). Using MP-BGP to distribute the reachability of MAC addresses over a MPLS network brings the same operational control and scale of L3VPN to L2VPN.

The EVPN solution provides a common base for all Ethernet service types including E-LAN, E-Line, E-Tree (including E-Tree* from TR-221), E-Access and enables these services to be created such that they can span across domains. In addition to the common base above, BGP MPLS-based EVPNs also provide solutions for the requirements in RFC 7209 [37] including:

- Multi-homing: with all active forwarding and load balancing from CE to CE. VPLS can only support multi-homing with single active mode.
- Flow based load balancing and multipath
- Multicast optimization: must be able to support P2MP MPLS LSPs and MP2MP MPLS LSPs. VPLS only supports P2MP MPLS LSPs.
- Fast convergence to minimize downtime and packet loss.
- Support MAC mobility to support cloud services.

6 Reference Architecture

6.1 General Reference Architecture

Figure 1 provides a generic overview of how Carrier Ethernet Services can be deployed using a BGP MPLS-based EVPN infrastructure, including basic reference points and their functional roles. Depending on the application, non-MEF-defined Ethernet Attachment Circuits and Attachment Circuits providing User-to-Network interfaces complying with Metro Ethernet Forum definitions (MEF UNI) are supported. Multi-domain connectivity and external handoff are also supported.

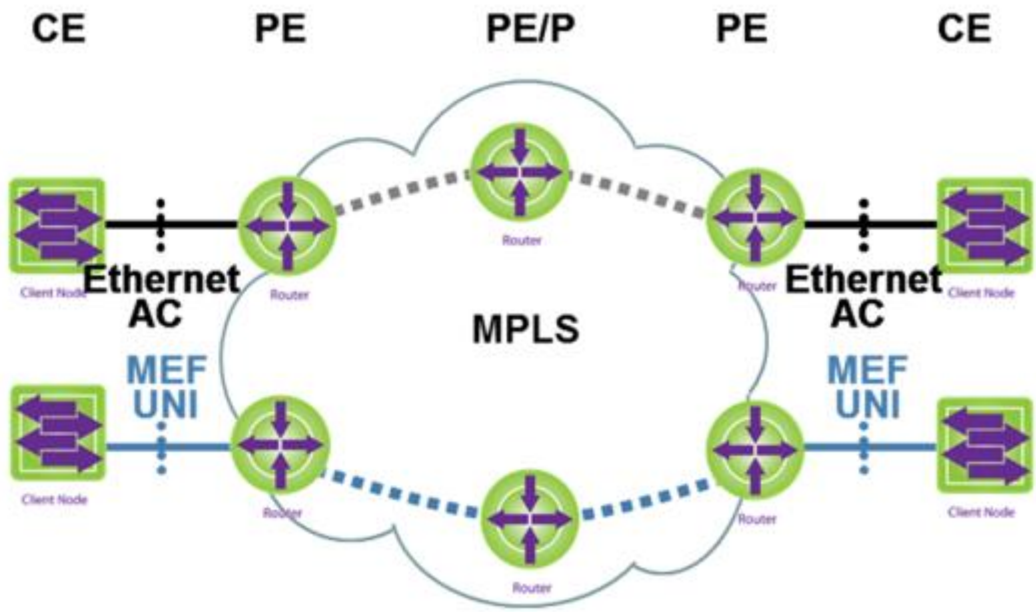

Figure 1 Reference Architecture

Defined as business interfaces supporting the service handoff between different parties (between user and provider or between providers, respectively), UNI has two functions:

1. provide reference points for network demarcation
2. provide associated functionality

For deploying Metro Ethernet Forum-compliant Ethernet Services over MPLS, PE nodes need to support the corresponding MEF UNI functionality at Attachment Circuit interfaces.

6.2 MPLS for Carrier Ethernet in Broadband Access & Aggregation

6.2.1 Multi-Service Broadband Access & Aggregation

For Multi-Service Broadband access and aggregation architecture see Section 6.2.1/TR-224 [4].

6.2.2 TR-178 Architectures

There are two reference architectures that are being used to represent TR-178 networks: 1) MPLS-enabled access and 2) TR-101. For architectural details of MPLS-enabled access nodes and TR-101 see Section 6.2.2/TR-224 [4].

7 Signaling and Routing

This section specifies the signaling protocol used to establish the underlying MPLS tunnel. Traffic engineered PSN tunnels must be used when specific path (e.g. for protection purpose), QoS or bandwidth constraints are required.

7.1 LSP Signaling

One of the following provisioning and signaling procedures are used for LSPs.

[R-1] PE and P routers supporting MPLS TE and non-TE LSPs MUST support one or both of the following methods:

- Static provisioning
- Dynamic signaling

[R-2] Both of the following methods MUST be supported by PE and P routers for dynamically signaled PSN tunnel LSPs:

- LDP is used to set up, maintain and release LSP tunnels per RFC 5036 [25].
- RSVP-TE is used to set up, maintain and release LSPs for traffic engineered tunnels per RFC 3209 [10] and RFC 5151 [27]. When traffic engineering is needed on the LSP, RSVP-TE MUST be used.

[R-3] When co-routed bidirectional LSPs are required, GMPLS-RSVP-TE as per RFC 3473 [12] MAY be supported by PE and P routers.

7.1.1 Multi-area LSP Signaling

Several operators have multi-area networks for scalability. Link state Interior Gateway Protocols (IGPs) such as OSPF (RFC 2328 [9]) and IS-IS (RFC 1195 [7]) allow dividing networks into areas or levels so as to increase routing scalability within a routing domain.

Further some operators' L2VPN networks span different geographical areas. To support these networks, it is necessary to support inter-area and inter-AS (Autonomous System) Multiprotocol Label Switching (MPLS) LSPs.

An "MPLS Domain" is considered to be any collection of network elements implementing MPLS within a common realm of address space or path computation responsibility. Examples of such domains include Autonomous Systems, Interior Gateway Protocol (IGP) routing areas, and GMPLS overlay networks.

Signaling extensions for inter-area LSPs (that is, LSPs that traverse at least two IGP areas) are required to ensure MPLS connectivity between PEs located in distinct IGP areas.

7.1.1.1 Multi-area RSVP-TE Signaling

Inter-domain TE LSPs can be supported by one of three options as specified in RFC 5151 [27] and given below:

- contiguous LSPs
- nested LSPs
- stitched LSPs.

Note: While the specifications below cover both MPLS Traffic Engineering (TE) and GMPLS, the intent is to use the TE coverage.

Contiguous

A contiguous TE LSP is a single TE LSP that is set up across multiple domains using RSVP-TE signaling procedures described in Section 7.1.

Nested

One or more TE LSPs may be nested within another TE LSP as described in RFC 4206 [20]. This technique can be used to nest one or more inter-domain TE LSPs into an intra-domain hierarchical LSP (H-LSP). The label stacking construct is used to achieve nesting in packet networks.

To improve scalability, it may be useful to aggregate LSPs by creating a hierarchy of such LSPs.

[R-4] PE routers SHOULD support establishment of RSVP-TE LSPs using LSP hierarchy as per RFC 4206 [20].

Stitched

LSP stitching signaling procedures are described in RFC 5150 [26]. This technique can be used to stitch together shorter LSPs (LSP segments) to create a single, longer LSP. The LSP segments of an inter-domain LSP may be intra-domain LSPs or inter-domain LSPs.

The process of stitching LSP segments results in a single, end-to-end contiguous LSP in the data plane. But in the control plane, each segment is signaled as a separate LSP (with distinct RSVP sessions) and the end-to-end LSP is signaled as yet another LSP with its own RSVP session. Thus, the control plane operation for LSP stitching is very similar to that for nesting.

[R-5] PE routers SHOULD support establishment of RSVP-TE LSPs using LSP stitching as per RFC 5150 [26].

7.1.1.2 Multi-area LDP Signaling

RFC 5283 [28] facilitates the establishment of Label Switched Paths (LSPs) that would span multiple IGP areas in a given Autonomous System (AS).

[R-6] PE routers SHOULD support establishment of inter-area LSPs using LDP as per RFC 5283 [28].

7.2 Routing

[R-7] One or both of the following methods MUST be supported by PE and P routers:

- Static routing
- Dynamic routing

[R-8] Both of the following methods MUST be supported by PE and P routers to exchange routing information to facilitate dynamic LSP signaling:

- OSPF (RFC 2328 [9])
- IS-IS (RFC 1195 [7])

[R-9] Traffic engineering extensions of OSPF and IS-IS are used to exchange traffic attributes for RSVP-TE tunnels. If TE is supported, both of the following methods MUST be supported by PE and P routers:

- OSPF-TE (RFC 3630 [16])
- IS-IS-TE (RFC 5305 [30])

8 OAM

OAM in Carrier Ethernet networks was developed to provide fault management and performance monitoring tools for network links and end-to-end EVCs. Figure 2 below shows the components of OAM when the MEF services are provided using EVPN.

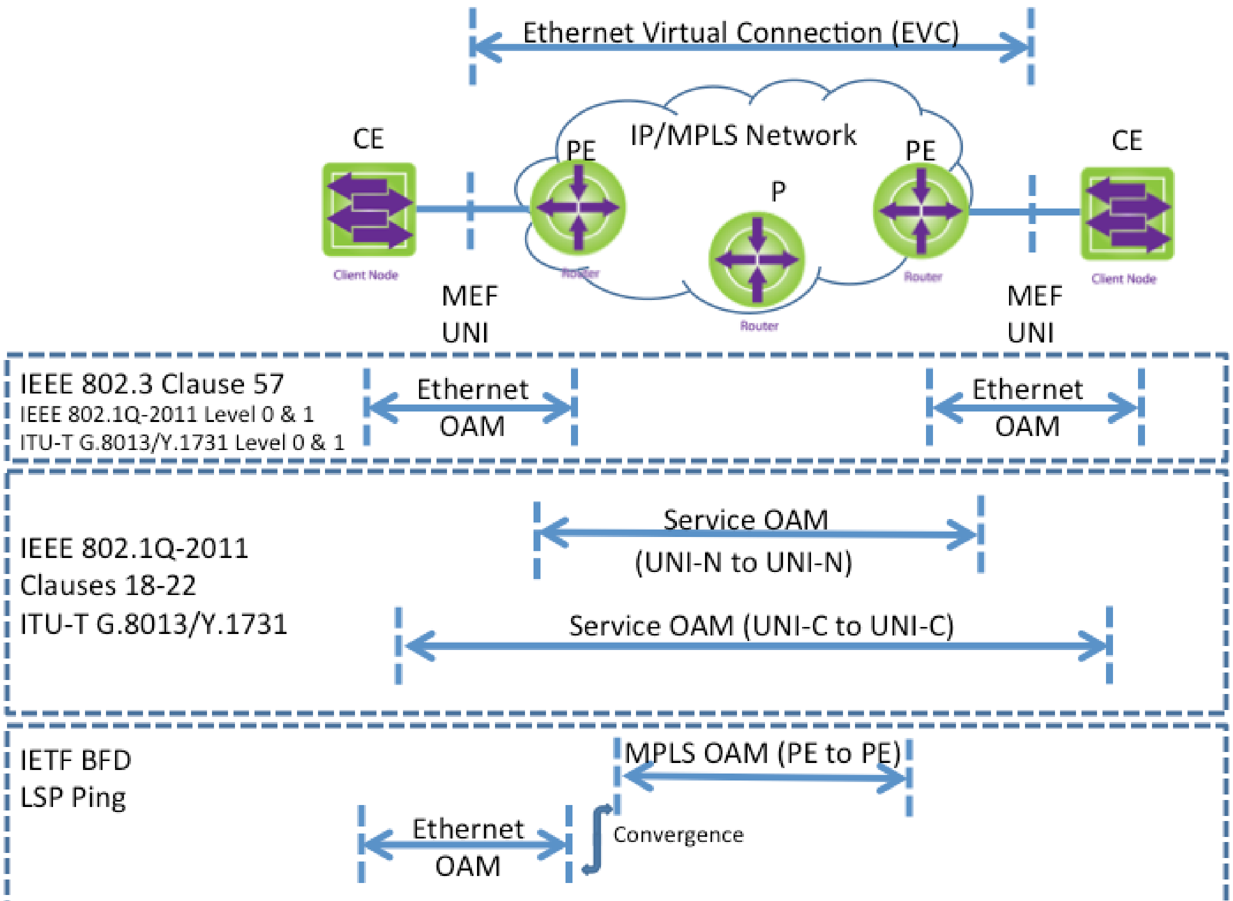

Figure 2 Components of OAM

8.1 Ethernet OAM

The OAM functions for the UNI should use the Ethernet OAM as defined in Link OAM (Clause 57/IEEE 802.3 [5]) and/or ITU-T G.8013/Y.1731 [40].

8.1.1 Link OAM

The PE supports Ethernet Link OAM, when the user is directly connected to the network demarcation point (i.e., at the MEF UNI in Figure 2). Link OAM provides OAM functions for network access segments (UNI-C to UNI-N). Link OAM provides for Ethernet Link Fault Detection, Monitoring and Loopback for access links.

[R-10] The PE MUST support link OAM Active mode as per Clause 57.2.9.1/IEEE 802.3 [5].

[R-11] The PE MUST support initiating OAM Discovery process as per Subclause 57.3.2.1/IEEE 802.3 [5].

[R-12] The PE MUST support sending informational OAM Protocol Data Units (OAMPDU) as per Subclause 57.2.10/IEEE 802.3 [5].

[R-13] The PE MUST support sending Event Notification OAMPDUs as per Subclause 57.2.10/IEEE 802.3 [5].

[R-14] The PE MUST support sending loopback control OAMPDUs as per Subclause 5.2.11/IEEE 802.3 [5].

[R-15] The PE MAY support sending Organization specific OAMPDU per Subclause 57/IEEE 802.3 [5].

[R-16] The PE MAY support sending Variable Request OAMPDUs as per Subclause 57.4.3.3/IEEE 802.3 [5].

8.1.2 ITU-T G.8013/Y.1731

The OAM functions defined in ITU-T G.8013/Y.1731 can be used for OAM of the UNI between the CE and the PE.

[R-17] The PE MUST support sending and receiving OAM frames at level 0 (as recommended in G.8013/Y.1731 [40]).

8.2 MEF Service OAM

The Carrier Ethernet Services are provided between one User Network Interface (UNI) and one or more UNIs. A network operator must be able to manage the services using Service OAM (SOAM). The network operator's service OAM is originated at the PE's UNI-N.

[R-18] The PE MUST support SOAM at the EVC SOAM level 4 as described in MEF 30 [43].

[R-19] OAM frames, sent at SOAM levels 5, 6, or 7, as described in MEF 30, are sent as user data and MUST be carried transparently.

[R-20] The PE MAY support sending and receiving SOAM frames across the UNI at the UNI SOAM level 1, as described in MEF 30 [43].

See Section 12.4.8 for information on performance monitoring.

8.3 MPLS OAM

This section describes techniques to perform OAM for the underlying MPLS tunnels used to support Ethernet services. OAM is an important and fundamental functionality in an MPLS network. OAM contributes to the reduction of operational complexity, by allowing for efficient and automatic detection, localization, handling and diagnosis of defects. OAM functions, in general, are used for fault-management, performance monitoring, and protection-switching applications.

8.3.1 LSP OAM

This section describes techniques to perform OAM for the underlying MPLS LSPs used in an EVPN application.

LSP-Ping and Bidirectional Forwarding Detection (BFD) RFC 5880 [32] are OAM mechanisms for MPLS LSPs RFC 5884 [34]. Further it is desirable that the OAM traffic is sent in-band in an LSP. The following OAM mechanisms are supported:

[R-21] The PE MAY support GAL and G-ACh per LSP, as per RFC 5586 [31].

8.3.1.1 BFD for MPLS LSPs

BFD monitors the integrity of the LSP for any loss of continuity defect. In particular, it can be used to detect a data plane failure in the forwarding path of an MPLS LSP.

[R-22] PE and P routers MUST support BFD for MPLS LSPs as per RFC 5884 [34].

8.3.1.2 Detecting MPLS Data Plane Failures

LSP Ping is used to perform on-demand Connectivity Verification (CV) and Route Tracing functions. It provides two modes: “ping” mode and “traceroute” mode.
In "ping" mode (basic connectivity check), the packet should reach the end of the path, at which point it is sent to the control plane of the egress LSR, which then verifies whether it is indeed an egress for the FEC.

[R-23] PE and P routers MUST support “ping” mode as per RFC 4379 [22].

RFC 6424 [35] enhances the mechanism for performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels and when LSP stitching [RFC5150] is in use.

[R-24] PE and P routers MUST support enhanced MPLS ping and traceroute as per RFC 6424 [35].

In "traceroute" mode (fault isolation), the packet is sent to the control plane of each transit LSR, which performs various checks that it is indeed a transit LSR for this path; this LSR also returns further information that helps check the control plane against the data plane.

[R-25] PE and P routers SHOULD support “traceroute” mode as per RFC 4379 [22].

The LSP Ping Reply modes as defined in Section 3/RFC 4379 [22] apply as shown in Table 1.

Reply Mode	Echo request	Echo Reply
Reply via an IPv4/IPv6 UDP packet (code value 2)	MUST	MUST
Reply via application level control channel (code value 4)	MAY	MAY

Table 1 LSP Ping Reply Modes

The following subsections of Section 3.2/RFC 4379 [22] concerning Target FEC Stack apply as follows:

[R-26] When LDP is supported - LDP IPv4 prefix as defined in Section 3.2.1/RFC 4379 [22] MUST be supported.

[R-27] When RSVP is supported - RSVP IPv4 LSP as defined in Section 3.2.3/RFC 4379 [22] MUST be supported.

[R-28] When BGP is supported – BGP-labeled IPv4 prefix as defined in Section 3.2.11/RFC 4379 [22] MUST be supported.

[R-29] When LDP is supported - LDP IPv6 prefix as defined in Section 3.2.2/RFC 4379 [22] SHOULD be supported.

[R-30] When RSVP is supported - RSVP IPv6 LSP as defined in Section 3.2.4/RFC 4379 [22] SHOULD be supported.

[R-31] When BGP is supported – BGP-labeled IPv6 prefix as defined in Section 3.2.12/RFC 4379 [22] SHOULD be supported.

8.3.2 Convergence

This section specifies the requirements for the recovery mechanisms from PE to CE network (AC link) failures. The recovery procedures are described in Section 17/RFC 7432 [39].

[R-32] The PE routers MUST support recovery from PE to CE network failures (AC link failures) as per Section 17.3/RFC 7432 [39].

[R-33] The PE routers MUST support recovery from PE failures as per Section 17.2/RFC 7432 [39].

9 QoS

The MPLS network supporting the Carrier Ethernet services has to provide QoS and service level agreements. The QoS capabilities must be end-to-end, which includes both ACs and MPLS domains. Usually an MPLS network will support guaranteeing sufficient bandwidth if available to support new and existing Carrier Ethernet connections conforming to all SLA metrics including protection mechanisms.

DiffServ-TE classes of service is used to support MEF 23.1 [45] "CoS Labels" to achieve a particular level of performance. MPLS DiffServ-TE enables the advantages of both DiffServ and TE. The DiffServ-TE requirement is to make separate bandwidth reservations for different classes of traffic. RFC 3564 [14] provides the concept of a class type (CT).

The following capabilities are to be supported by the PEs:

[R-34] The PE MUST support at least 4 CoS and associated service metrics (e.g. delay, delay variation, packet loss) as defined in MEF 22.1 [44] "EVC Requirements".

[R-35] The PE SHOULD support Connection Admission Control to guarantee sufficient bandwidth is available to support new connection conforming to all SLA metrics defined in MEF 10.2 [42].

[R-36] The PE SHOULD support Differentiated Service aware MPLS traffic engineering as per RFC 4124 [19].

[R-37] The ingress PE MUST map the PCP (in the PRI field of the 802.1Q VLAN tag IEEE 802.1Q [6]) into the TC field of the MPLS label stack.

9.1 Tunnel CoS mapping and marking

Two types of LSPs are defined in RFC 3270 [11]:

[R-38] The PE and P routers MUST support E-LSP as per Section 1.2/RFC 3270 [11]: LSPs which can transport multiple Ordered Aggregates, so that the TC field of the MPLS Shim Header conveys to the LSR the PHB to be applied to the packet (covering both information about the packet's scheduling treatment and its drop precedence).

[R-39] The PE and P routers MAY support L-LSP as per Section 1.3/RFC 3270 [11]: LSPs which only transport a single Ordered Aggregate, so that the packet's scheduling treatment is inferred by the LSR exclusively from the packet's label value while the packet's drop precedence is conveyed in the TC field of the MPLS Shim Header.

[R-40] The PE MUST support COS marking in the TC bits of the LSP labels.

[R-41] The PE MUST support the Pipe model as per RFC 3270 [11].

10 PSN resiliency

In EVPN, the PEs are connected over an underlying PSN infrastructure. For EVPN resiliency, the PSN must necessarily be resilient. When the PEs are connected by an MPLS infrastructure then the resiliency mechanisms in MPLS such as fast reroute (FRR) are required. This section lists the resiliency requirements for MPLS. If the MPLS infrastructure is run over another layer (e.g. a L1 network) the resiliency requirements of the other layers are considered to be outside the scope of this section. When resiliency mechanisms are available at multiple layers the resiliency mechanism at a layer must be triggered only after a sufficient delay to let the resiliency mechanism of the underlying layer to take effect.

MPLS resiliency requires failure detection mechanisms and LSP recovery mechanisms. To speed up total recovery time, local-repair mechanisms with pre-computed, pre-established alternate/backup paths should be used whenever possible. Section 10.1 lists the failure detection requirements and Section 10.2 lists the requirements for LSP recovery.

MPLS resiliency is also affected by the restart of control-plane protocols. The MPLS requirements to support resiliency of control protocols are listed in Section 10.3.

10.1 Failure detection

The failure detection mechanism that triggers the recovery mechanisms should have a low failure detection time and also a low overhead. In order for the deployment to allow a choice of routing protocols, the failure detection mechanism should be independent of specific routing protocols. The Bidirectional Forwarding Detection (BFD) protocol specified in RFC 5880 [32] provides such a mechanism. For MPLS LSP BFD requirements see Section 8.3.1.

[R-42] The PE and P routers MUST support BFD for single hops as per RFC 5881 [33]

10.2 LSP recovery

The LSP recovery mechanism should support local repair mechanisms with pre-computed and pre-established alternate/backup paths for both RSVP-TE RFC 3209 [10] and LDP RFC 5036 [25] signaled LSPs. Recovery from different types of failure such as link, node, etc. should be supported.

[R-43] The PE and P routers MUST support the facility backup method of doing fast reroute (FRR) for RSVP-TE LSP Tunnels as per RFC 4090 [18].

[R-44] The PE and P routers SHOULD support the one-to-one backup method of doing fast reroute (FRR) for RSVP-TE LSP Tunnels as per RFC 4090 [18].

[R-45] The PE and P routers MUST support the loop-free alternates (LFA) method of FRR for LDP LSPs as per RFC 5286 [29] as well as support LFA FRR for the IGP on whose routes LDP depends.

10.3 Control plane resiliency

To prevent LSPs from going down due to control-plane protocols restart, the graceful restart control-plane resiliency mechanism is required.

[R-46] The PE and P routers MUST support RSVP-TE graceful restart as specified in Section 9/RFC 3473 [12] as well as graceful restart for the routing protocols on which RSVP-TE path computation depends.

[R-47] The PE and P routers MUST support LDP graceful restart as specified in RFC 3478 [13] as well as graceful restart for the routing protocols on whose routes LDP depends.

[R-48] The PE and P routers SHOULD support OSPF graceful restart as specified in RFC 3623 [15].

[R-49] The PE and P routers SHOULD support IS-IS graceful restart as specified in RFC 5306 [17].

11 BGP MPLS-Based Ethernet VPN

This section covers the generic BGP MPLS-based Ethernet VPN requirements. Specific requirements such as multicast that are applicable to a subset of Ethernet VPN services (e.g. EP-LAN, EVP-LAN, etc) are covered in subsequent sections.

EVPN overcomes the limitations of current E-Line and E-LAN services supported by VPLS (RFC 4761 [23], RFC 4762[24]) and VPWS. EVPN provides flexible multi-homing with all-active redundancy mode, MAC learning using control plane, multicast optimization, provisioning simplicity and network resiliency between edge nodes.

The EVPN specification supports several ways for PE nodes to connect, but this TR only supports use of MPLS, which enable easy interworking with TR-224 [4] based Ethernet services.

11.1 Reference Architecture and Overview

Figure 3 describes EVPN for the next-generation of Ethernet services. An EVPN instance comprises of CEs connected to PEs, which are part of the MPLS network. The PEs provide virtual Layer 2 bridge connectivity between CEs. The PEs are connected by an underlying MPLS network that provides QoS and resiliency.

Unlike VPLS, which uses only data-plane based MAC learning, EVPN uses the control plane based MAC learning for remote MACs. EVPN uses MP-BGP to distribute MAC routes and allows fine-grained control over MAC route distribution.

EVPN instance (EVI) is an EVPN routing and forwarding instance on a PE. If a CE is multi-homed to two or more PEs, the set of Ethernet links constitute an Ethernet Segment (ES). Each Ethernet Segment is identified using a unique Ethernet Segment identifier (ESI). For additional details see RFC 7432 [39].

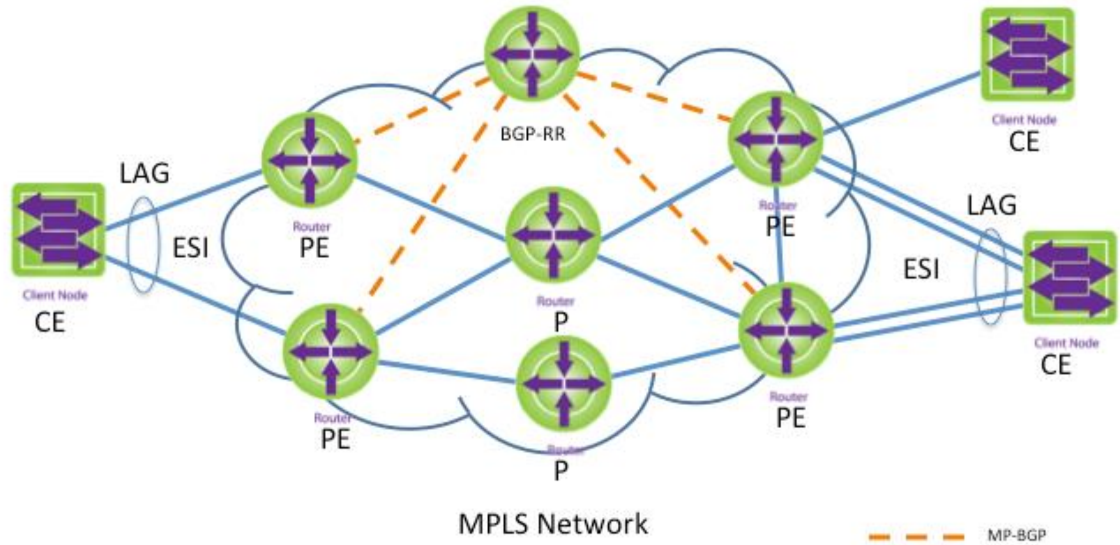

Figure 3 EVPN Architecture for Ethernet services using BGP MPLS

11.2 EVPN Service Interfaces

EVPN defines several types of service interfaces. See Section 6/RFC 7432 [39] for details. These service interfaces are consistent with MEF-defined services and provide easy migration to EVPN infrastructure for even richer service offerings. The various types of service interfaces include mapping of specific VLANs or their bundles and even allow service awareness when they are mapped to EVPN instances. The requirements for the support of various service interfaces are specified in the subsections of the respective services.

11.2.1 VLAN-Based Service Interfaces

This service interface supports a single broadcast domain or VLAN per-EVPN instance. This service interface can be used to support E-LAN or E-Line for a single broadcast domain with customer VLANs having local significance. Ethernet frames transported over an MPLS network remain tagged with the originating VID. VID translation can be performed on the destination PE.

11.2.2 VLAN Bundle Service Interfaces

This service interface supports a bundle of VLAN over one EVPN instance. Multiple VLANs share the same bridge. It supports an N:1 mapping between VLAN ID and MAC-VRF. This service interface requires that MAC addresses are unique across VLANs of the EVI and VID translation is not allowed.

This service interface also supports a special case known as a port-based VLAN Bundle service interface, where all the VLANs on a port are part of the same service and map to the same bundle.

11.2.3 VLAN-Aware Bundle Service Interfaces

This service interface is an additional service interface defined in EVPN that is not supported by TR-224 or VPLS. It provides customers with a single E-LAN service for multiple broadcast domains. With this service interface, an EVPN instance consists of multiple broadcast domains or VLANs, with each VLAN having its own bridge domain. Like the VLAN bundle service interface, this interface supports N:1 mapping between VLAN ID and EVI. Since bridge domains are separate, it allows for local VID translation.

This service interface also supports a special case known as a port-based VLAN-Aware bundle service interface, where all the VLANs on a port are part of the same service and map to the same bundle.

11.3 Data Plane

11.3.1 Underlying PSN transport

[R-50] The PEs MUST support MPLS as the underlying PSN transport as specified in Section 4/RFC 7432 [39].

11.3.2 VPN encapsulation

To distinguish packets received over the PSN destined to different EVPN instances, MPLS labels must be used as described in Section 4/RFC 7432 [39]. The specific data plane operations applicable to a service are specified in the subsections of the respective service.

11.3.3 VID Translation

[R-51] The PEs MUST support VID translation for packets received from the PSN and sent to the CE, when supporting service interfaces as specified in Section 6/RFC 7432 [39].

11.3.4 Frame Ordering

Section 18/RFC 7432 specifies frame ordering. In order to avoid misordering, it is recommended that P routers not use deep packet inspection to do ECMP.

[R-52] The P routers SHOULD NOT do deep packet inspection for ECMP. RFC 6790 [36] specifies techniques so that P routers do effective load balancing without the need for deep packet inspection.

11.4 Control Plane

The EVPN PEs signal and learn MAC address over the control plane. RFC 7432 adds BGP extended communities, which allow PE routers to advertise and learn MAC addresses and Ethernet segments. This is one of the major differences with the VPLS solution, which relies on data-plane learning. EVPN added four Route types and communities. For additional details see RFC 7432 [39].

[R-53] The PEs MUST support MP-BGP as a control protocol for EVPN as specified in Section 4 and 7/RFC 7432 [39].

Note: The detailed control protocol requirements of MP-BGP are specified in the subsections of the respective services.

With MPLS data plane, BGP routes also signal the MPLS labels associated with MAC addresses and Ethernet segments. This separates EVPN from a VPLS solution. EVPNs do not use Pseudowires.

11.5 Multi-Homing and Load balancing

Due to rapid increase of data traffic, running the network in active/standby mode can be inefficient. In addition to better link utilization, multi-homed connections also offer greater resiliency and reliability against the failure of one connection or node. Multi-homing includes the ability of establishing multiple connections between PEs and to load-balance across those connections. For additional details on multi-homing see Section 8/RFC 7432 [39]. EVPNs supports both single-active and all-active multi-homing with load balancing. VPLS only supports single-active multi-homing.

With support for both all-active per-service and all-active per-flow multi-homing, EVPNs enables better load balancing across peering PEs as compared to VPLS that cannot load balance across peering PEs.

It must be possible to connect a CE to two or more PEs for purposes of multi-homing and load balancing as specified in Section 8 and 14/RFC 7432 [39].

11.5.1 All-Active Redundancy Mode

All-active redundancy mode allows the CE device to connect via a "single" Ethernet bundle to multiple PEs using LAG. All the PEs must be allowed to forward traffic to/from that Ethernet Segment.

[R-54] PE router MUST support "All-Active redundancy mode" as specified in Section 14/RFC 7432 [39].

11.5.2 Single-Active Redundancy Mode

In this mode, when a CE is connected to two or more PEs over an Ethernet segment, only a single PE must be allowed to forward traffic to/from that Ethernet Segment. In this mode the CE device connect via "separate" Ethernet bundles to multiple PEs.

[R-55] PE router MUST support "Single-Active redundancy mode" as specified in Section 14/RFC 7432 [39].

11.6 Fast Convergence

Section 17/RFC 7432 provides failure recovery from different types of network failures. VPLS relies on the underlying MPLS capabilities such as Fast Reroute. Lack of all-active multi-homing in VPLS makes it difficult to achieve fast restoration in case of an edge node or edge link failure.

[R-56] The PEs MUST support convergence as specified in Section 17/RFC 7432 [39].

12 EVPN enabled multipoint to multipoint Ethernet VPN services (E-LAN)

The EVPN technology enables the creation of multipoint-to-multipoint Ethernet VPN services over an MPLS network. EVPN can be used to create the EP-LAN and EVP-LAN services of the E-LAN service type defined by MEF 6.2. A high-level reference architecture of how these services are architected using EVPN along with the list of the supported service attributes is described in Section 12.1 and 12.2. In addition to the Carrier Ethernet defined service characteristics, EVPN significantly enhances important service characteristics such as reliability and scalability. The EVPN requirements for multipoint-to-multipoint Ethernet VPN services are listed in Section 12.3.

12.1 Ethernet Private LAN (EP-LAN)

The Ethernet Private LAN (EP-LAN) service uses a multipoint-to-multipoint EVC. In a multipoint EVC, two or more UNIs must be associated with one another. The EP-LAN service is defined to provide CE-VLAN tag preservation and tunneling of key Layer 2 control protocols. A key advantage of this service is that VLANs can be configured across the sites without any need to coordinate with the service provider.

EP-LAN provides connectivity to customers with multiple sites, such that all sites appear to be on the same local area network. Each interface is configured for "All to One Bundling". EP-LAN supports CE-VLAN CoS preservation. Service multiplexing is disabled on the UNI.

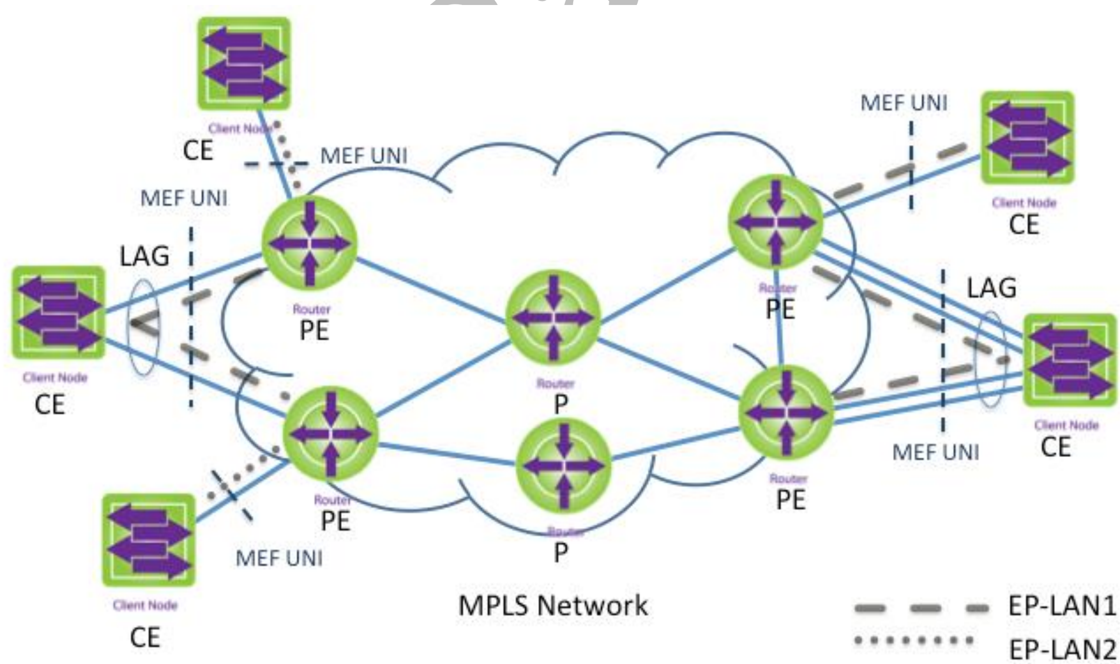

Figure 4 Ethernet Private LAN (EP-LAN) Service

12.2 Ethernet Virtual Private LAN (EVP-LAN)

The Ethernet Virtual Private LAN (EVP-LAN) service allows service multiplexing at the UNI. It allows users of an E-LAN service type to interconnect their UNIs and at the same time access other services (e.g. E-Line). Figure 5 shows an example of multiple services access from a single UNI. In this example, the user has an EVP-LAN service for multipoint data connectivity and an EVPL service (P2P EVC) for accessing a value-add service from one of the UNIs.

Bundling can be used on the UNI in the EVP-LAN service and supports CE-VLAN tag preservation. All to One Bundling is disabled.

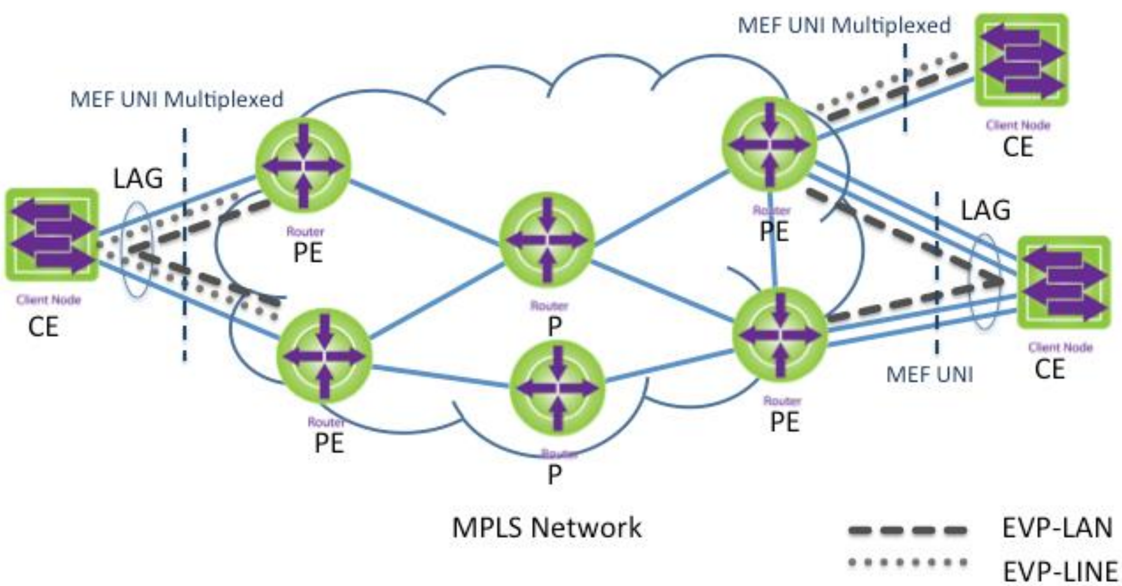

Figure 5 Ethernet Virtual Private LAN (EVP-LAN) Service

12.3 EVPN for establishing EP-LAN and EVP-LAN

Section 5 provides an overview of PPVPN in MPLS networks. It also outlines the comparison of Layer 2 Ethernet VPNs in MPLS networks using VPLS and EVPNs.

EVPN provides support for the E-LAN service type in MPLS networks as described in Section 11. RFC 7432 describes procedures for BGP MPLS-based Ethernet VPNS. EVPN requires extensions to existing IP/MPLS protocols. EVPN supports both provisioning and signaling for Ethernet VPNs and incorporate flexibility for service delivery over Layer 3 networks.

The PE supports BGP MPLS-based Ethernet VPN signaling and provisioning as specified in [R-53] of Section 11.4.

12.3.1 Service Interfaces

EVPN supports several service connectivity options for delivering MEF services. They are provided in Section 11.2. MEF service requirements for EP-LAN and EVP-LAN are different. For EP-LAN, each interface is configured for "All to One Bundling". For EVP-LAN "All to One Bundling" is disabled. VLAN-Aware Bundle service interface is not supported in VPLS-based implementation (TR-224). When interworking with TR-224, it is recommended not to use VLAN-Aware Bundle service interface. This service interface provides a more flexible service offering and is commonly used for data center interconnect.

12.3.1.1 Service Interfaces for EP-LAN

[R-57] The PE routers MUST support Port-based Service Interface as defined in Section 6.2.1/RFC 7432 [39].

[R-58] The PE routers SHOULD support Port-Based VLAN-Aware Service Interface as defined in Section 6.3.1/RFC 7432 [39].

12.3.1.1.1 Data Center Considerations

If the PE router is designed for use in a data center interconnect environment, the following requirement is applicable:

[R-59] The PE routers MUST support Port-Based VLAN-Aware Service Interface as defined in Section 6.3.1/RFC 7432 [39].

12.3.1.2 Service Interfaces for EVP-LAN

[R-60] The PE routers MUST Support VLAN-based Service Interface as defined in Section 6.1/RFC 7432 [39].

[R-61] The PE routers MUST support VLAN Bundle Service Interface as defined in Section 6.2/RFC 7432 [39]

[R-62] The PE routers SHOULD support VLAN-Aware Bundle Service Interface as defined in Section 6.3/RFC 7432 [39].

12.3.1.2.1 Data Center Considerations

If the PE router is designed for use in a data center interconnect environment, the following requirement is applicable:

[R-63] The PE routers MUST support VLAN-Aware Bundle Service Interface as defined in Section 6.3/RFC 7432 [39].

12.3.2 Data plane

The requirements for data plane per Section 11.3 are applicable.

12.3.2.1 Local learning

[R-64] The PE MUST be able to do data-plane learning of MAC addresses using IEEE Ethernet learning procedures for packets received from the CEs connected to it as specified in Section 9.1/RFC 7432 [39].

12.3.2.2 Remote learning

[R-65] The PE MUST be able to do control-plane learning of MAC addresses using MP-BGP's MAC Advertisement route for CEs that are connected to remote PEs as specified in Section 9.2/RFC 7432 [39].

12.3.3 Tunnel signaling

The PEs are connected by MPLS Label Switch Paths (LSPs) acting as PSN tunnels. Traffic Engineered PSN tunnels must be used when specific path (e.g. for protection purpose), QoS, or bandwidth constraints are required.

[R-66] PE and P routers MUST support dynamic signaling to setup both TE LSPs and routed LSPs. See Section 7.1 for details.

12.3.4 Routing

The requirements for routing per Section 7.2 are applicable.

12.3.5 Multi-Homing and Load balancing

The requirements for multi-homing and load balancing per Section 11.5 are applicable.

12.3.5.1 Load balancing at intermediate nodes

When load balancing, packets that belong to a given 'flow' must be mapped to the same port. Intermediate P nodes have no information about the type of the payload inside the LSP. Intermediate LSR should make a forwarding choice based on the MPLS label stack. In order to avoid any mis-ordering of frames, the requirements specified in Section 11.3.4 apply.

The PE that has knowledge of the Ethernet service (e.g. Bundling or multiclass service) can take further action. IETF RFC 6790 [36] provide methods of assigning labels to flows, or flow groups, such that Label Switching Routers can achieve better load balancing.

[R-67] The PE SHOULD support Entropy Labels as per RFC 6790 [36].

12.3.6 OAM

12.3.6.1 Ethernet Link OAM

The Ethernet link OAM is supported as per Section 8.1.1.

12.3.6.2 Label Switched Paths (LSP) OAM

LSP OAM is supported as per Section 8.3.1.

12.3.6.3 MEF Service OAM

MEF service is supported as per Section 8.1.2.

12.3.7 Convergence

Failure recovery from different types of network failure is supported as per Section 8.3.2.

12.3.8 PSN Resiliency

PSN resiliency is supported as per Section 10.

12.3.8.1 Fast Convergence

Fast convergence is supported as per Section 11.6.

12.3.9 Multicast and Broadcast

[R-68] PE routers SHOULD support multicast and broadcast traffic as per Section 16/RFC 7432 [39].

12.3.10 QoS

In general, an E-LAN service type can provide a best effort service with no performance assurance. In certain cases, an E-LAN service type can be defined with performance objectives (see Section 9.2/MEF 6.2 [50].

[R-69] PE routers SHOULD support the QoS mapping as per Section 9.

12.3.11 Security

[R-70] PE routers MUST support security as per Section 19/RFC 7432 [39].

12.4 Support of service attributes for EP-LAN and EVP-LAN

Section 9.2/MEF 6.2 [50] specifies the E-LAN service type that is the basis for LAN services. Section 10.3 and 10.4/MEF 6.2 [50] provides service attributes and parameters for EP-LAN and EVP-LAN services respectively. TR-224 refers to MEF 6.1 and MEF 10.2. TR-350 uses the backward compatible subset of the revised specifications MEF 6.2 [50] (see Appendix B) and MEF 10.3 [49] to achieve the equivalent function.

Some of the service attributes and parameters are provided by Ethernet physical interface and service provisioning (e.g., Physical medium, Speed, Mode, MAC layer, EVC type, maximum number of EVCs, etc.). This section only describes those service attributes and parameters that are relevant to transporting the EVPN traffic over PSN.

12.4.1 Bandwidth Profile

A bandwidth profile defines how rate enforcement of Ethernet frames is applied at an UNI. Bandwidth profiles enable offering service bandwidth below the UNI access speed (aka Speed) and limit the amount of traffic entering the network per the terms of the SLA.

For LAN services, bandwidth profiles can be optionally specified per-UNI (ingress and egress), per-EVC (ingress and egress), and/or per CoS (ingress and egress). An E-LAN service can be provided as a best effort service without any bandwidth guarantee.

[R-71] A PE SHOULD support the bandwidth profile algorithm as per the portion of Section 12/MEF 10.3 [49] that is backward compatible with MEF 10.2 [42].

In order to support bandwidth profile, technique such as admission control and Diffserv-TE as specified in Section 9 are used.

12.4.2 Bundling

Section 9.12/MEF 10.3 [49] specifies the bundling service attribute. Bundling implies “A UNI attribute in which more than one CE-VLAN ID can be associated with an EVC”. All to one bundling enabled is a special case of bundling. It implies “A UNI attribute in which all CE-VLAN IDs are associated with a single EVC”. Table 12/MEF 10.3 [49] provides valid combinations for “All to one bundling” and “Service multiplexing” attributes.

EP-LAN must have “All to one bundling” attribute enabled. For EVP-LAN the bundling attribute can be enabled or disabled. However, for EVP-LAN “All to one bundling” must be disabled.

12.4.3 CE-VLAN ID preservation for EVC

CE-VLAN ID preservation service attribute defines whether the CE-VLAN ID is preserved (unmodified) across the EVC.

For EP-LAN, CE-VLAN ID preservation must be enabled and CE-VLAN ID is preserved for EVC over the PSN.

For EVP-LAN, if CE-VLAN ID preservation is enabled, CE-VLAN ID is preserved for EVC over the PSN.

Note: If CE-VLAN ID preservation is enabled, No VID translation is supported for the EVC. EVP-LAN can support VID translation when using EVPN service type "VLAN-Based Service type" and "VLAN-Aware Bundle service interface".

12.4.4 CE-VLAN CoS preservation for EVC

CE-VLAN CoS preservation service attribute defines whether the CE-VLAN CoS bits are preserved (unmodified) across the EVC.

For EP-LAN, CE-VLAN CoS preservation must be enabled (see Table 15/MEF 6.2 [50]) and CE-VLAN CoS is preserved for EVC over PSN.

For EVP-LAN, CE-VLAN CoS preservation can be either enabled or disabled (see Table 18/MEF 6.2). In an EVC with CE-VLAN CoS preservation is enabled, the EVPN preserves the CoS bits over PSN.

12.4.5 EVC Maximum Service Frame Size

The mapping from "EVC Maximum Service Frame Size" to "EVC MTU" is provided in Table 40 Appendix B of MEF 6.2 [50].

The EVC Maximum Service Frame Size size is configurable with a default value of 1600 byte.

When Ethernet frames are transported in MPLS networks, the MPLS packet includes the labels, and the EVC frame as payload. The path MTU is the largest packet size that can traverse this path without fragmentation. The ingress PE can use Path MTU Discovery to find the actual path MTU.

[R-72] PE SHOULD support configurable EVC Maximum Service Frame Size of at least 1600 bytes (see Table 6/MEF 6.2).

12.4.6 Frame delivery

The frame delivery policy rules enable the service provider to specify how different frame types are handled by a PE. They enable setting specific rules for forwarding, discarding or conditionally forwarding specific frame types. The frame types used by the rules are:

- Unicast
- Multicast
- Broadcast

[R-73] PE MUST support setting policy function of frame delivery rules for forwarding, discarding or conditionally forwarding unicast, multicast and broadcast frames per EP-LAN and EVP-LAN services.

12.4.7 Layer 2 control protocols

The Layer 2 control protocol processing is independent of the EVC at the UNI. L2CP handling rules are set according to the definition of Section 8/MEF 6.1.1 [48] and differ per-service type. The PE policy function supports setting of rules for handling L2CP per service type.

EVC L2CP handling per service type can be set to:
- Discard – Drop the frame.
- Peer can be applicable: For example L2CP/LAMP, Link OAM, Port Authentication, and E-LMI.
- Tunnel – Pass to the egress UNI.

[R-74] PE MUST support policy function setting of rules for handing L2CP per service type as specified in Section 8/MEF 6.1.1 [48].

Note: This specification only supports MEF 6.1.1 for L2CP processing requirements. Support of multiple-CEN L2CP MEF 45 [51] is outside the scope of this document.

12.4.8 EVC performance

The performance parameters indicate the quality of service for that service instance. They consist of the following:
- Availability
- Frame Delay
- Frame delay variation
- Frame loss ratio

The requirements for support of CoS and mapping are specified in QoS Section 9.

[R-75] The PE MUST support MEF SOAM performance monitoring as per MEF 35 [47].

For transport of SOAM see Section 8.2.

13 EVPN enabled point to point Ethernet VPN services (E-Line)

The Ethernet Line service is defined in MEF 6.2 [50] and is based on a point-to-point Ethernet Virtual Connection (EVC). Virtual Private Wire Service (VPWS) is a Layer 2 VPN service used to emulate the E-Line service type in an MPLS network. draft-ietf-bess-evpn-vpws **Error! Reference source not found.** describes how EVPN can be used to support VPWS in IP/MPLS networks. The use of EVPN mechanisms for VPWS brings the benefits of EVPN to point-to-point services. EVPN can be used to create the Ethernet Private Line (EPL) and Ethernet Virtual Private Line (EVPL) services of the E-Line service type defined by MEF 6.2 [50].

A high-level reference architecture of how these services are architected using EVPN along with the list of the supported service attributes is described in Section 13.1 and 13.2. In addition to the Carrier Ethernet-defined service characteristics, EVPN significantly enhances important service characteristics such as reliability and scalability. The EVPN-VPWS requirements for point-to-point Ethernet VPN services are listed in Section 13.3.

13.1 Ethernet Private Line (EPL)

Ethernet Private Line (EPL), uses a point-to-point EVC between two UNIs to provide a high degree of transparency for service frames between UNIs it interconnects. The service frames, headers, and most Layer 2 protocols are identical at both the source and destination UNI. It does not allow for service multiplexing; that is, a dedicated UNI (physical interface) is used for the EPL. The figure below shows EPL service.

A key advantage of this service is that VLANs can be configured across the sites without any need to coordinate with the service provider. Each interface is configured for “All to One Bundling”. EPL supports CE-VLAN CoS preservation. For cases where an ingress bandwidth profile is applied, the CE is expected to shape traffic.

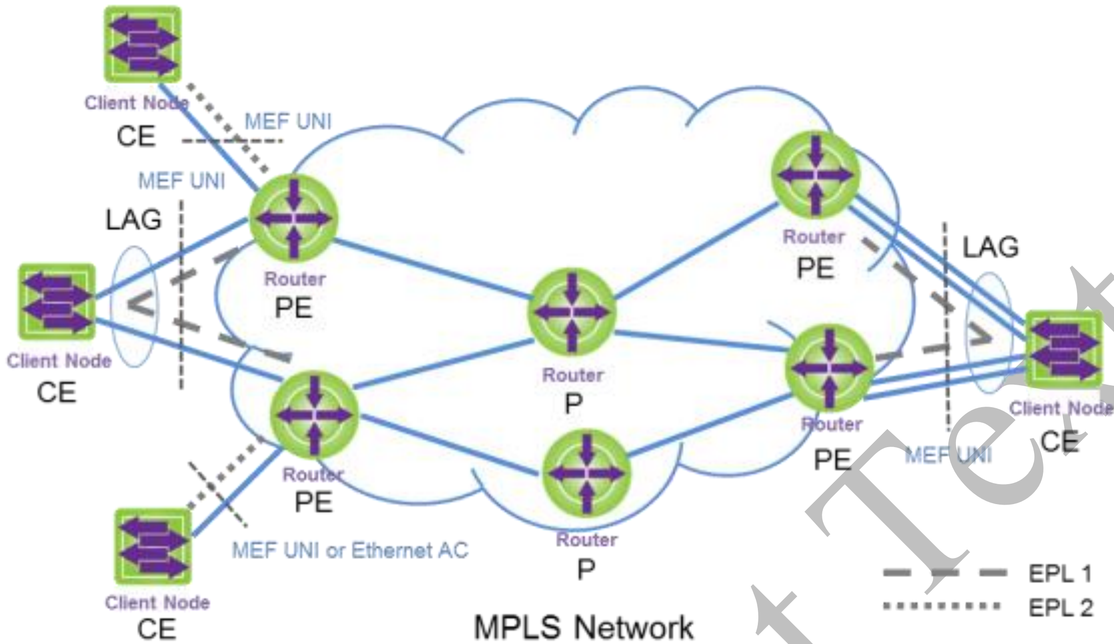

Figure 6 Ethernet Private Line (EPL) Service

Ethernet Private Line (EPL):

- It replaces a TDM private line
- Dedicated UNIs for Point-to-Point connections
- Single Ethernet Virtual Connection (EVC) per UNI
- The most popular Ethernet service type due to its simplicity

13.2 Ethernet Virtual Private Line (EVPL)

Ethernet Virtual Private Line (EVPL) uses a point-to-point EVC between two UNIs, but does not provide full transparency as with the EPL service. The EVPL service also allows for service multiplexing, which means that more than one EVC can be supported at the UNI. Because service multiplexing is permitted, some service frames may be sent to one EVC, while other service frames may be sent to other EVCs. The service definition for EVPL is specified in Section 10.1/MEF 6.2 [50].

EVPL is commonly used for connecting subscriber hub and branch locations. It allows users of an E-Line service type to interconnect their UNIs and at the same time access other services (e.g. E-LAN). Figure 7 shows an example of multiple services access from a single UNI. In this example, the user has an EVP-LAN service for multipoint data connectivity and an EVPL service (P2P EVC) for accessing a value-add service from one of the UNIs. Bundling can be used on the UNI in the EVPL service and supports CE-VLAN tag preservation. All to One Bundling is disabled.

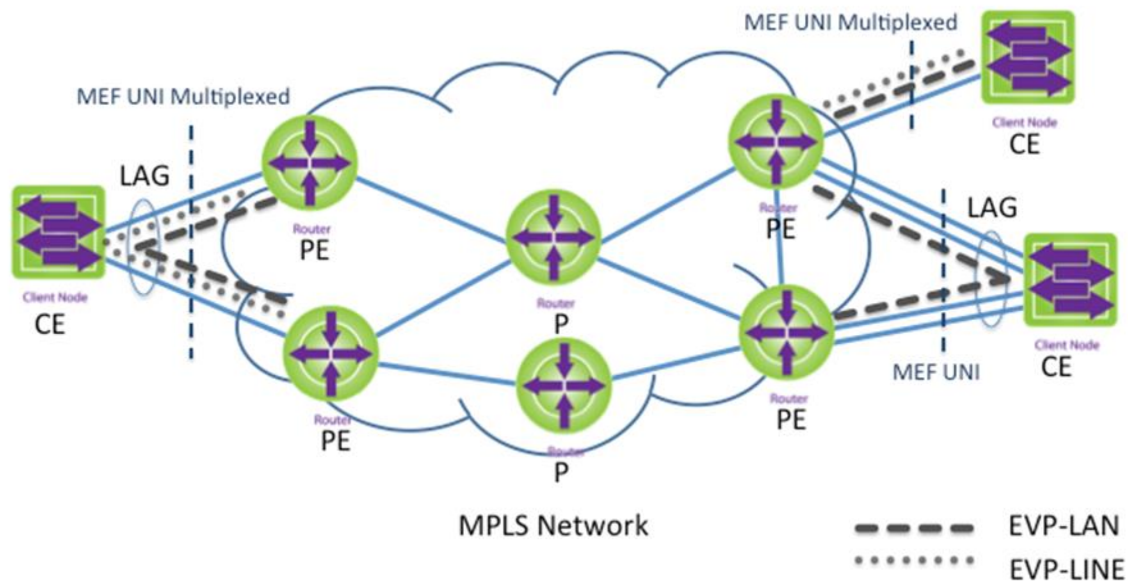

Figure 7 Ethernet Virtual Private Line (EVPL) Service

13.3 EVPN for establishing EPL and EVPL

VPWS support in EVPN draft-ietf-bess-evpn-vpws **Error! Reference source not found.** describes how EVPN can be used to support VPWS in IP/MPLS networks. EVPN enables the following service capabilities: single-active as well as all-active multi-homing; flow-based load balancing; and policy. It eliminates the need for PWs for point-to-point Ethernet services. EVPN has the ability to forward customer traffic at the PE without any MAC lookup because the MPLS label associated with the EVI is used. EVPL can be considered as a VPWS with only two-attachment circuits (AC) or MEF-UNIs. For additional details see draft-ietf-bess-evpn-vpws **Error! Reference source not found.**.

RFC 7432 [39] describes procedures for BGP MPLS-based Ethernet VPNs. EVPN requires extensions to existing IP/MPLS protocols. EVPN supports both provisioning and signaling for Ethernet VPNs and incorporates flexibility for service delivery over Layer 3 networks.

The PE supports BGP MPLS-based Ethernet VPN signaling and provisioning as specified in requirement [R-53] of Section 11.4.

13.3.1 Service Interfaces

EVPN-VPWS supports several service connectivity options for delivering MEF services. They are provided in Section 2/draft-ietf-bess-evpn-vpws **Error! Reference source not found.**. This service does not use “VLAN-Aware Bundle Service interface”. .

MEF service requirements for EPL and EVPL are different. For EPL, each interface is configured for “All to One Bundling”. For EVPL “All to One Bundling” is disabled.

13.3.1.1 Service Interfaces for EPL

[R-76] The PE routers MUST support Port-based Service Interface as defined in Section 6.2.1/RFC 7432 [39].

13.3.1.2 Service Interfaces for EVPL

[R-77] The PE routers MUST Support VLAN-based Service Interface as defined in Section 2.1/draft-ietf-bess-evpn-vpws **Error! Reference source not found.**.

[R-78] The PE routers MUST support VLAN Bundle Service Interface as defined in Section 2.2/draft-ietf-bess-evpn-vpws **Error! Reference source not found.**.

13.3.2 Operation

PE routers must support point to point deployments as defined in Section 4/draft-ietf-bess-evpn-vpws **Error! Reference source not found.**.

13.3.3 Data Plane

The requirements for data plane per Section 11.3.1 and 11.3.2 are applicable.

13.3.3.1 Local learning

The requirements for local learning per Section 12.3.2.1 are applicable.

13.3.3.2 Remote learning

The requirements for remote learning per Section 12.3.2.2 are applicable.

13.3.4 BGP extensions

EVPN-VPWS uses the Ethernet A-D per-EVI route to signal VPWS services. An EVPN instance MUST NOT be configured with both VPWS service instance and EVPN multipoint services. draft-ietf-bess-evpn-vpws **Error! Reference source not found.** Section 3.1 adds a EVPN Layer 2 a ttributes extended community.

[R-79] The PE routers MUST support EVPN Layer 2 attributes extended community as defined in Section 3.1/draft-ietf-bess-evpn-vpws **Error! Reference source not found.**.

13.3.5 Tunnel signaling

The requirements for tunnel signaling per Section 12.3.3 are applicable.

13.3.6 Routing

The requirements for routing per Section 7.2 are applicable.

13.3.7 Multi-Homing and Load balancing

The requirements for multi-homing and load balancing per Section 11.5 are applicable.

13.3.7.1 Load balancing at intermediate nodes

The requirements for load balancing at intermediate nodes per Section 12.3.5.1 are applicable.

13.3.8 OAM

13.3.8.1 Ethernet Link OAM

The Ethernet link OAM is supported as per Section 8.1.1.

13.3.8.2 Label Switched Paths (LSP) OAM

LSP OAM is supported as per Section 8.3.1.

13.3.8.3 MEF Service OAM

MEF service is supported as per Section 8.2.

13.3.9 Convergence

Failure recovery from different types of network failure is supported as per Section 8.3.2.

13.3.10 PSN Resiliency

PSN resiliency is supported as per Section 10.

13.3.10.1 Fast Convergence

Fast convergence is supported as per Section 11.6.

13.3.11 QoS

E-Line service type can be defined with performance objectives (see Section 9.1/MEF 6.2 [50]). PE routers MUST support the QoS mapping as per Section 9.

13.3.12 Security

[R-80] PE routers MUST support security as per Section 19/RFC 7432 [39].

13.4 Support of service attributes for EPL and EVPL

Section 9.1/MEF 6.2 [50] specifies the Ethernet Line (E-Line) Service Type that is the basis for a broad range of Point-to-Point services. Section 10.1 and 10.2/MEF 6.2 [50] provides service definitions for the Ethernet Private Line (EPL) Service and the Ethernet Virtual Private Line (EVPL) Service that are based on the E-Line Service Type by defining the required Service Attributes and related parameter values. The service attributes are tabulated in Section 8/MEF 6.2 [50].

TR-224 refers to MEF 6.1 [41] and MEF 10.2 [42]. TR-350 uses the backward compatible subset of the revised specifications MEF 6.2 [50] and MEF 10.3 [49] to achieve the equivalent function. This addresses interworking with TR-224.

This section only describes those Service Attributes and parameters that are relevant to transporting the service over the PSN.

13.4.1 Bandwidth Profile

A bandwidth profile characterizes the lengths and arrival times of Service frames at a reference point. Bandwidth profiles enable offering service bandwidth below the UNI access speed (aka Speed) and limit the amount of traffic entering the network per the terms of the SLA.

For the EPL and EVPL service, bandwidth profiles can be optionally specified per-UNI (ingress and egress), per-EVC (ingress and egress), and/or per CoS (ingress and egress) as specified in Section 10.1 and 10.2/MEF 6.2 [50]. The EPL and EVPL service can be provided as a best effort service without any bandwidth guarantee.

Note: Bandwidth profiles can be optionally provisioned per UNI (ingress and egress). They are not carried by EVPN protocols.

[R-81] A PE SHOULD support the bandwidth profile algorithm as per the portion of Section 12/MEF 10.3 [49] that is backward compatible with MEF 10.2 [42].

In order to support bandwidth profile, technique such as admission control and Diffserv-TE as specified in Section 9 are used.

13.4.2 Bundling

Bundling is defined in Section 2/MEF 10.3 [49] as "A UNI Service attribute in which more than one CE-VLAN ID can be associated with an EVC" and is described in Section 9.12/MEF 10.3 [49]. All to one bundling defined as "A UNI attribute in which all CE-VLAN IDs are associated with a single EVC" is a special case of Bundling Enabled and is described in Section 9.13/MEF 10.3 [49]. Table 12/MEF 10.3 [49] provides valid combinations for "Service multiplexing", "Bundling" and "All to one bundling" attributes.

EPL must have the "Bundling" attribute as disabled and "All to one bundling" attribute as enabled. For EVPL the "Bundling" attribute can be enabled or disabled, but "All to one bundling" must be disabled.

13.4.3 CE-VLAN ID preservation for EVC

For EPL, CE-VLAN ID preservation must be enabled and CE-VLAN ID is preserved for EVC over the PSN.

For EVPL, if CE-VLAN ID preservation is enabled, CE-VLAN ID is preserved for EVC over the PSN.

Note: If CE-VLAN ID preservation is enabled, No VID translation is supported for the EVC. EVPL can support VID translation when using EVPN service type "VLAN-Based Service type".

13.4.4 CE-VLAN CoS preservation for EVC

CE-VLAN CoS preservation service attribute defines whether the CE-VLAN CoS bits are preserved (unmodified) across the EVC.

For EPL, CE-VLAN CoS preservation may be enabled or disabled (see Table 12/MEF 6.2 [50]). For EVPL CE-VLAN CoS is supported as per sec 10.3/MEF 6.2 [50] and may be enabled or disabled. In an EVC with CE-VLAN CoS preservation is enabled, the EVPN preserves the CoS bits over PSN.

13.4.5 EVC Maximum Service Frame Size

The mapping from "EVC Maximum Service Frame Size" to "EVC MTU" is provided in Table 40 Appendix B of MEF 6.2 [50].

The EVC Maximum Service Frame Size is configurable with a default value of 1600 bytes.

When Ethernet frames are transported in MPLS networks, the MPLS packet includes the labels, and the EVC frame as payload. The path MTU is the largest packet size that can traverse this path without fragmentation. The ingress PE can use Path MTU Discovery to find the actual path MTU.

[R-82] PE SHOULD support configurable EVC Maximum Service Frame Size of at least 1600 bytes (see Table 6/MEF 6.2 [50]).

13.4.6 Frame delivery

The frame delivery policy rules enable the service provider to specify how different frame types are handled by a PE. They enable setting specific rules for forwarding, discarding or conditionally forwarding specific frame types. The frame types used by the rules are:

- Unicast
- Multicast
- Broadcast

[R-83] PE MUST support setting policy function of frame delivery rules for forwarding, discarding or conditionally forwarding unicast, multicast and broadcast frames per EPL and EVPL services.

13.4.7 Layer 2 control protocols

The Layer 2 control protocol processing is independent of the EVC at the UNI. L2CP handling rules are set according to the definition of Section 8/MEF 6.1.1 [48] and differ per-service type. The PE policy function supports setting of rules for handling L2CP per service type.

EVC L2CP handling per service type can be set to:

- Discard – Drop the frame.
- Peer can be applicable: For example L2CP/LAMP, Link OAM, Port Authentication, and E-LMI.
- Tunnel – Pass to the egress UNI.

[R-84] PE MUST support policy function setting of rules for handing L2CP per service type as specified in Section 8/MEF 6.1.1 [48].

Note: This specification only supports MEF 6.1.1 for L2CP processing equirements. Support of multiple-CEN L2CP MEF 45 [51] is outside the scope of this document.

13.4.8 EVC performance

The performance parameters indicate the quality of service for that service instance. They consist of the following:

- Availability
- Frame Delay
- Frame delay variation
- Frame loss ratio

The requirements for support of CoS and mapping are specified in QoS Section 9.

[R-85] The PE MUST support MEF SOAM performance monitoring as per MEF 35 [47].

For transport of SOAM see Section 8.2.

13.4.9 Multiple Class of Service

MEF 6.1 [41] introduced services that were allowed multiple Classes of Service. The different ways that EVCs can be created to support such a service is defined in Section 11.4.9/MEF 6.1 [41].

[R-86] The PE MUST support a single class of service for an EVC by using a single value for the TC bits in the MPLS label stack.

[R-87] The PE MUST support a multiple class of service for an EVC by using multiple values for the TC bits in the MPLS label stack.

14 EVPN enabled Rooted-multipoint Ethernet VPN services (E-Tree)

MEF Ethernet Tree (E-Tree) service type is based upon a Rooted-multipoint Ethernet Virtual connection. The EVPN technology enables the creation of E-Tree services over an MPLS network. EVPN can be used to create the EP-Tree and EVP-Tree services of the E-Tree service type defined by MEF 6.2 [50]. In an E-Tree service, endpoints are labeled as either Root or Leaf sites. Root sites can communicate with all other sites. Leaf sites can communicate with Root sites but not with other Leaf sites.

This service could be useful for internet access or video over IP application, such as multicast/broadcast packet video.

High-level services are described in Section 14.1and 14.2. In addition to the Carrier Ethernet defined service characteristics, EVPN significantly enhances important service characteristics such as reliability and scalability. The EVPN requirements for E-Tree Ethernet VPN services are listed in Section 14.3.

14.1 Ethernet Private Tree (EP-Tree)

The Ethernet Private Tree (EP-Tree) service uses a Rooted-multipoint EVC. In a multipoint EVC, two or more UNIs must be associated with one another. The EP-Tree service is defined to provide CE-VLAN tag preservation and tunneling of key Layer 2 control protocols. A key advantage of this service is that VLANs can be configured across the sites without any need to coordinate with the service provider.

Each interface is configured for “All to One Bundling”. EP-LAN supports CE-VLAN CoS preservation. Service multiplexing is disabled on the UNI.

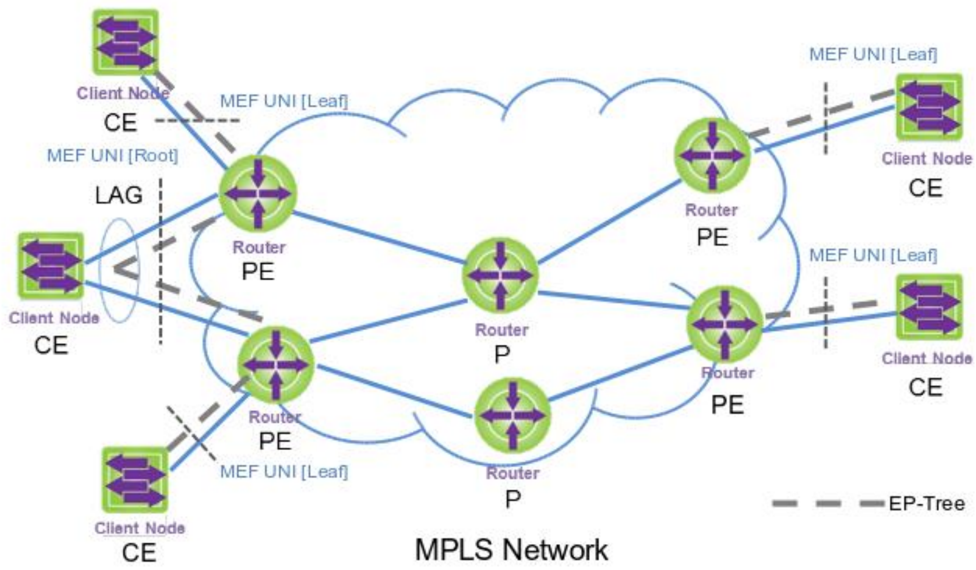

Figure 8 Ethernet Private Tree (EP-Tree) Service

14.2 Ethernet Virtual Private Tree (EVP-Tree)

The Ethernet Virtual Private Tree (EVP-Tree) service allows service multiplexing at the UNI. It allows users of an E-Tree service to interconnect their UNIs and at the same time access other services (e.g. E-Line). Figure 9 shows an example of multiple services access from a single UNI. In this example, the user has an EVP-Tree service for Rooted-multipoint data connectivity and an EVPL service (P2P EVC) for accessing a value-add service from one of the UNIs.

Bundling can be used on the UNI in the EVP-Tree service and supports CE-VLAN tag preservation. All to One Bundling is disabled.

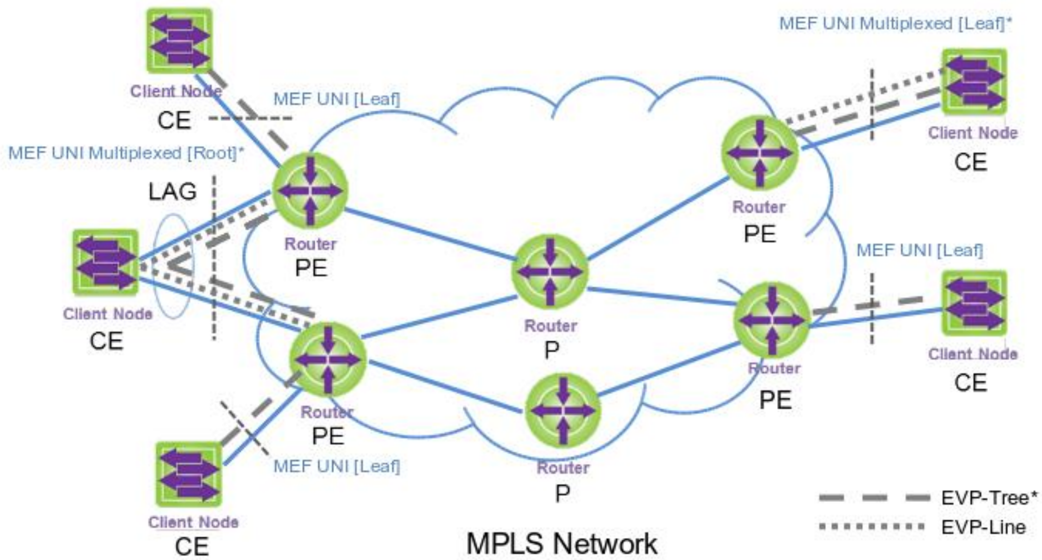

Figure 9 Ethernet Virtual Private Tree (EVP-Tree) Service

14.3 EVPN for establishing EP-Tree and EVP-Tree

L2VPN service provides multipoint to multipoint connectivity for Ethernet across an IP or MPLS Network. A generic E-LAN/E-Tree service is always bidirectional in the sense that ingress frames can originate at any endpoint in the service.

The only difference between E-LAN and E-Tree is:

- E-LAN has Root endpoints only, which implies there is no communication restriction between endpoints.
- E-Tree has both Root and Leaf endpoints, which implies there is a need to enforce communication restriction between Leaf endpoints.

RFC 7387[38] proposes the solution framework for supporting E-Tree service in MPLS networks. The document identifies additional components to emulate E-Tree service using Ethernet LAN service over MPLS network.

EVPN provides support for the E-LAN service type in MPLS networks as described in Section 11. RFC 7432 [39] describes procedures for BGP MPLS-based Ethernet VPNS. EVPN requires extensions to existing IP/MPLS protocols. EVPN supports both provisioning and signaling for Ethernet VPNs and incorporates flexibility for service delivery over Layer 3 networks.

The PE supports BGP MPLS-based Ethernet VPN signaling and provisioning as specified in [R-53] of Section 11.4.

draft-ietf-bess-evpn-etree [52] describes how EVPN can be used to support E-Tree in IP/MPLS networks. The RFC discusses how the functional requirements for E-Tree service can be implemented efficiently with EVPN.

14.3.1 E-Tree support in EVPN

14.3.1.1 E-Tree Scenarios and EVPN Support

The support for E-Tree services in EVPN is divided into different scenarios, depending on the site association (Root/Leaf) per PE or Access Ethernet Segment. They are described in section 2/ draft-ietf-bess-evpn-etree [52].

[R-88] The PE routers MUST Support Leaf or Root site(s) per PE as defined in Section 2.1/ draft-ietf-bess-evpn-etree [52].

[R-89] The PE routers MUST Support Leaf or Root site(s) per Attachment Circuit (AC) as defined in Section 2.2/ draft-ietf-bess-evpn-etree [52].

14.3.1.2 E-Tree Extended Community

A new BGP extended community is used for leaf indication. They are described draft-ietf-bess-evpn-etree [52].

[R-90] The PE routers MUST Support BGP extended community encoding as defined in Section 5/ draft-ietf-bess-evpn-etree [52].

14.3.1.3 E-Tree Operation for EVPN

RFC 7432 [39] has inherent capability to support E-Tree service. It adds a new BGP extended community for leaf indication. The E-Tree operations for EVPN are described in section 3/ draft-ietf-bess-evpn-etree [52].

[R-91] The PE routers MUST Support E-Tree operation for EVPN as defined in Section 3/ draft-ietf-bess-evpn-etree [52].

14.3.2 Service Interfaces

EVPN supports several service connectivity options for delivering MEF services. They are provided in Section 11.2. MEF service requirements for EP-LAN and EVP-LAN are different. For EP-LAN, each interface is configured for “All to One Bundling”. For EVP-LAN “All to One Bundling” is disabled.

14.3.2.1 Service Interfaces for EP-Tree

[R-92] The PE routers MUST support Port-based Service Interface as defined in Section 6.2.1/RFC 7432 [39].

14.3.2.2 Service Interfaces for EVP-Tree

[R-93] The PE routers MUST Support VLAN-based Service Interface as defined in Section 6.1/RFC 7432 [39].

[R-94] The PE routers MUST support VLAN Bundle Service Interface as defined in Section 6.2/RFC 7432 [39]

14.3.3 Data plane

The requirements for data plane per Section 12.3.2 are applicable.

14.3.4 Tunnel signaling

The requirements for tunnel signaling per Section 12.3.3 are applicable.

14.3.5 Routing

The requirements for routing per Section 7.2 are applicable.

14.3.6 Multi-Homing and Load balancing

The requirements for multi-homing and load balancing per Section 12.3.5 are applicable.

14.3.7 OAM

The requirements for OAM per Section 12.3.6 are applicable.

14.3.8 Convergence

Failure recovery from different types of network failure is supported as per Section 8.3.2.

14.3.9 PSN Resiliency

PSN resiliency is supported as per Section 10.

14.3.9.1 Fast Convergence

Fast convergence is supported as per Section 11.6.

14.3.10 Multicast and Broadcast

[R-95] PE routers SHOULD support multicast and broadcast traffic as per Section 16/RFC 7432 [39].

14.3.11 QoS

In general, an E-Tree service type can provide a best effort service with no performance assurance. In certain cases, an E-Tree service type can be defined with performance objectives (see Section 9.2/MEF 6.2 [50].

[R-96] PE routers SHOULD support the QoS mapping as per Section 9.

14.3.12 Security

[R-97] PE routers MUST support security as per Section 19/RFC 7432 [39].

14.4 Support of service attributes for EP-Tree and EVP-Tree

Section 9.3/MEF 6.2 [50] specifies the E-Tree service type. Section 10.5 and 10.6/MEF 6.2 [50] provides service attributes and parameters for EP-Tree and EVP-Tree services respectively. TR-224 refers to MEF 6.1 [41] and MEF 10.2. TR-350 uses the backward compatible subset of the revised specifications MEF 6.2 [50] (see Appendix B) and MEF 10.3 [49] to achieve the equivalent function.

Some of the service attributes and parameters are provided by Ethernet physical interface and service provisioning (e.g., Physical medium, Speed, Mode, MAC layer, EVC type, maximum number of EVCs, etc.). This section only describes those service attributes and parameters that are relevant to transporting the EVPN traffic over PSN.

14.4.1 Bandwidth Profile

A bandwidth profile defines how rate enforcement of Ethernet frames is applied at an UNI. Bandwidth profiles enable offering service bandwidth below the UNI access speed (aka Speed) and limit the amount of traffic entering the network per the terms of the SLA.

For Tree services, bandwidth profiles can be optionally specified per-UNI (ingress and egress), per-EVC (ingress and egress), and/or per CoS (ingress and egress). An E-Tree service can be provided as a best effort service without any bandwidth guarantee.

Note: Bandwidth profiles can be optionally provisioned. They are not carried by EVPN protocols.

[R-98] A PE SHOULD support the bandwidth profile algorithm as per the portion of Section 12/MEF 10.3 [49] that is backward compatible with MEF 10.2 [42].

In order to support bandwidth profile, technique such as admission control and Diffserv-TE as specified in Section 9 are used.

14.4.2 Bundling

Section 9.12/MEF 10.3 [49] specifies the bundling service attribute. Bundling implies “A UNI attribute in which more than one CE-VLAN ID can be associated with an EVC”. All to one bundling enabled is a special case of bundling. It implies “A UNI attribute in which all CE-VLAN IDs are associated with a single EVC”. Table 12/MEF 10.3 [49] provides valid combinations for “All to one bundling” and “Service multiplexing” attributes.

EP-Tree must have “All to one bundling” attribute enabled. For EVP-Tree the bundling attribute can be enabled or disabled. However, for EVP-Tree “All to one bundling” must be disabled.

14.4.3 CE-VLAN ID preservation for EVC

CE-VLAN ID preservation service attribute defines whether the CE-VLAN ID is preserved (unmodified) across the EVC.

For EP-Tree, CE-VLAN ID preservation must be enabled and CE-VLAN ID is preserved for EVC over the PSN.

For EVP-Tree, if CE-VLAN ID preservation is enabled, CE-VLAN ID is preserved for EVC over the PSN.

Note: If CE-VLAN ID preservation is enabled, No VID translation is supported for the EVC. EVP-Tree can support VID translation when using EVPN service type “VLAN-Based Service type”.

14.4.4 CE-VLAN CoS preservation for EVC

CE-VLAN CoS preservation service attribute defines whether the CE-VLAN CoS bits are preserved (unmodified) across the EVC.

For EP-Tree, CE-VLAN CoS preservation must be enabled (see Table 21/MEF 6.2 [50]) and CE-VLAN CoS is preserved for EVC over PSN.

For EVP-Tree, CE-VLAN CoS preservation can be either enabled or disabled (see Table 24/MEF 6.2). In an EVC with CE-VLAN CoS preservation is enabled, the EVPN preserves the CoS bits over PSN.

14.4.5 EVC Maximum Service Frame Size

The requirements for EVC maximum service frame size per Section 12.4.5 are applicable.

14.4.6 Frame delivery

The frame delivery policy rules enable the service provider to specify how different frame types are handled by a PE. They enable setting specific rules for forwarding, discarding or conditionally forwarding specific frame types. The frame types used by the rules are:

- Unicast
- Multicast
- Broadcast

[R-99] PE MUST support setting policy function of frame delivery rules for forwarding, discarding or conditionally forwarding unicast, multicast and broadcast frames per EP-Tree and EVP-Tree services.

14.4.7 Layer 2 control protocols

The requirements for Layer 2 control protocols per Section 12.4.7 are applicable.

14.4.8 EVC performance

The requirements for EVC performance per Section 12.4.8 are applicable.

End of Broadband Forum Technical Report TR-350