

Network Working Group
Internet Draft
Category: Standard Track
Expires: October 2012

Raymond Key, Huawei
Simon Delord, Alcatel-Lucent
Frederic Jounay, Orange CH
Wim Henderickx, Alcatel-Lucent
Lucy Yong, Huawei
Lizhong Jin, ZTE

April 16, 2012

Extension to VPLS for E-Tree
draft-key-12vpn-vpls-etree-07

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on October 16, 2012.

Abstract

This document proposes a simple and effective approach to emulate E-Tree services over an MPLS network. [Section 4](#) presents a minimal extension to the current VPLS architecture defined in [[RFC4761](#)] and [[RFC4762](#)] to fulfil the specific E-Tree requirement: Leaf cannot communicate with Leaf. Backward compatibility issues are addressed also.

Table of Contents

1. Introduction.....	3
2. Terminology.....	4
3. Reference Model.....	5
4. Extension to VPLS for E-Tree.....	6
4.1. AC Type.....	6
4.2. Control Word.....	6
4.3. Additional Action in Data Forwarding.....	6
5. Backward Compatibility.....	8
5.1. AC Type.....	8
5.2. Control Word.....	8
5.3. Additional Action in Data Forwarding.....	8
5.3.1. Ingress PE Supporting the Extension but Egress PE Not.....	8
5.3.2. Egress PE Supporting the Extension but Ingress PE Not.....	9
6. Optional Enhancements for Leaf-only PE.....	9
6.1. No PW between Leaf-only PEs.....	10
6.2. Not Forward Frame from Leaf AC to Leaf-only PE.....	10
7. Hierarchical VPLS.....	10
7.1. Hierarchical LDP VPLS.....	10
7.1.1. Spoke Connectivity for Bridging-Capable Devices.....	10
7.1.2. Multi-domain VPLS Service.....	11
7.1.3. Spoke Connectivity for Non-Bridging Devices.....	11
7.1.4. Hierarchical VPLS Model using Ethernet Access Network.....	12
7.2. Hierarchical BGP VPLS.....	12
8. Compliance with Requirements.....	12
9. Security Considerations.....	12
10. IANA Considerations.....	12
11. Acknowledgements.....	12
12. References.....	13
12.1. Normative References.....	13
12.2. Informative References.....	13
Appendix A. Control Word Scenarios.....	14
Appendix B. Data Forwarding Scenarios.....	15
Authors' Addresses.....	26
Intellectual Property and Copyright Statements.....	26

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

1. Introduction

A specific type of multipoint Ethernet services has been defined by Metro Ethernet Forum (MEF), called E-Tree [[MEF6.1](#)]. At the difference of E-LAN where there is no communication restriction between endpoints, on E-Tree each endpoint is designated as either a Root or a Leaf. Root can communicate with all other endpoints on the E-Tree, however Leaf can only communicate with Roots but not Leafs.

[Draft VPLS ETree Req] provides the functional requirements for MEF E-Tree support in VPLS.

E-LAN services are currently emulated via the use of the VPLS architecture.

This document presents a minimal extension to the current VPLS architecture defined in [[RFC4761](#)] and [[RFC4762](#)], with the objective to provide a simple and effective approach to fulfil the additional requirements of E-Tree compared to E-LAN.

Backward compatibility issues are also addressed in this document. PEs supporting the extension specified in this document and PEs not supporting such extension can inter-operate and participate in the same VPLS instance.

This document does not intend to address efficient packet replication or bandwidth optimisation, but the approach presented here does not prohibit future enhancements on those aspects.

In this document, "current standard" refers to [[RFC4385](#)], [[RFC4447](#)], [[RFC4448](#)], [[RFC4761](#)] and [[RFC4762](#)].

2. Terminology

E-Tree

An Ethernet VPN in which each Root AC can communicate with every other AC, whereas Leaf ACs can only communicate with Root ACs. Each AC on an E-Tree construct is designated as either a Root AC or a Leaf AC. There can be multiple Root ACs and Leaf ACs per E-Tree construct.

Root AC

An ingress frame at a Root AC can be delivered to one or more of any of the other ACs in the E-Tree. Please note that this AC is bidirectional.

Leaf AC

Ingress frame at a Leaf AC can only be delivered to one or more Root ACs in the E-Tree. Ingress frame at a Leaf AC MUST NOT be delivered to any Leaf ACs in the E-Tree. Please note that this AC is bidirectional.

3. Reference Model

Figure 1 below describes a generic reference model where PE1, PE2 and PE3 need to establish an E-Tree construct between different Ethernet endpoints. Each PE has 2 Root ACs and 2 Leaf ACs connected to a VSI. These VSIs are then linked together via Ethernet PWs.

In most use cases, an E-Tree construct has only a few Root ACs but many Leaf ACs. There may be only Root ACs or only Leaf ACs on a PE.

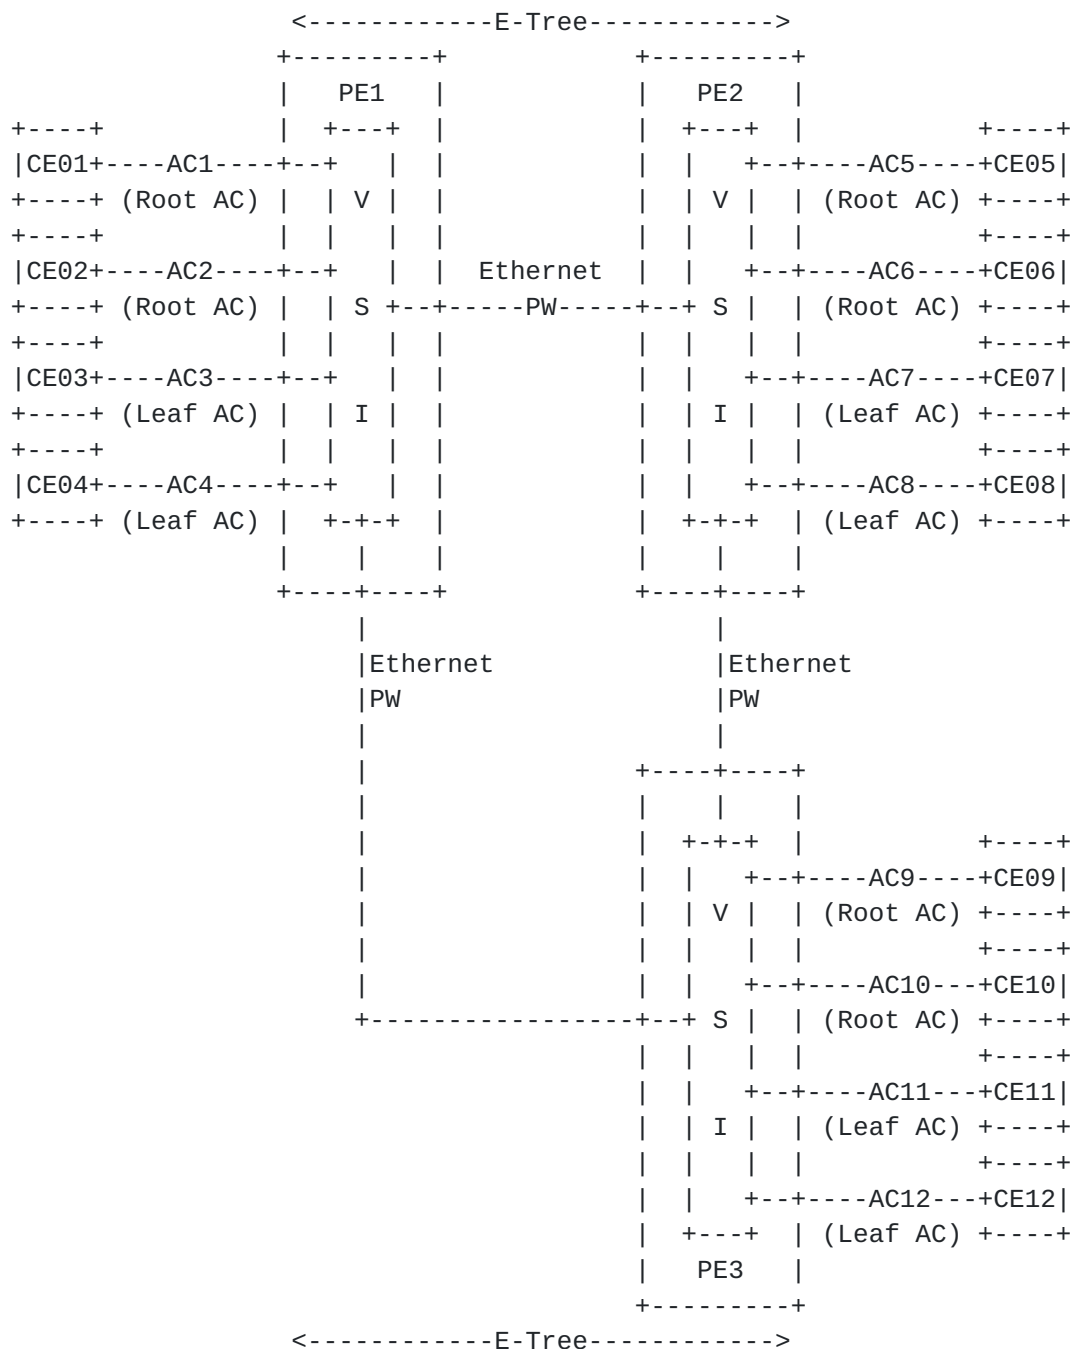


Figure 1: E-Tree Reference Model

Key, et al.

Expires October 2012

[Page 5]

With an E-Tree construct:

- A Root AC can receive from and transmit to any other ACs.
- A Leaf AC can receive from and transmit to any Root ACs.
- A Leaf AC cannot receive from and transmit to any other Leaf ACs.

This applies to all traffic, including Unicast Known, Unicast Unknown, Broadcast and Multicast.

When an Ethernet Frame is received on PE1 via AC1, the frame can be transmitted to any other local ACs on PE1 and via Ethernet PWs to any remote ACs on PE2 and PE3.

However when an Ethernet Frame is received on PE1 via AC3, the frame can be transmitted to any other local Root ACs on PE1 and via Ethernet PWs to any remote Root ACs on PE2 and PE3, but the frame cannot be transmitted to any local Leaf ACs on PE1 nor any remote Leaf ACs on PE2 and PE3.

4. Extension to VPLS for E-Tree

4.1 AC Type

Each AC connected to a specific VPLS instance on a PE MUST have an AC Type attribute, either Leaf AC or Root AC. The default value for AC Type attribute MUST be Root AC.

This AC Type is only locally configured on a PE and no signaling is required between PEs.

4.2 Control Word

A PE MUST be capable of sending and receiving the Control Word on Ethernet PW. Use of Control Word on Ethernet PW MUST be PREFERRED.

The procedure for negotiating the use of Control Word as per current standard is sufficient and MUST be followed. As a result, Control Word will always be used on Ethernet PW between two PEs when both PEs support the extension specified in this document.

Refer to [Appendix A](#) for different Control Word scenarios.

4.3 Additional Action in Data Forwarding

A PE MUST support the Control Word L-bit defined in [Draft CW L-bit].

A PE MUST perform the additional actions specified in Table 1 below.

The "Set CW L-bit" action and "Forward or Drop" decision are in addition to and performed after the following

- MAC-based forwarding decision as per current standard
- Loop free VPLS "split horizon" rule (MUST NOT forward traffic from one PW to another PW in the same VPLS mesh) as per current standard

		IF Conditions AND			THEN Actions	
Rule						
	Forward	Receive	CW	Set CW	Forward	
	Frame to	Frame from	L-bit	L-bit	or Drop	
1	Root AC	any AC/PW	any	n/a	Forward	
2	PW with no CW	any AC	n/a	n/a	Forward	
3	PW with CW	Root AC	n/a	Set to 0	Forward	
4	PW with CW	Leaf AC	n/a	Set to 1	Forward	
5	Leaf AC	Root AC	n/a	n/a	Forward	
6	Leaf AC	Leaf AC	n/a	n/a	Drop	
7	Leaf AC	PW with no CW	n/a	n/a	Forward	
8	Leaf AC	PW with CW	= 0	n/a	Forward	
9	Leaf AC	PW with CW	= 1	n/a	Drop	

Table 1: Additional Action in Data Forwarding

"Forward Frame to" is the result of MAC-based forwarding decision as per current standard.

"CW L-bit" is the bit 4 in the Ethernet PW Control Word as defined in [Draft CW L-bit] (the 5th bit in Control Word since the first bit is bit 0). The CW L-bit = 1 if and only if the payload Ethernet frame is from a Leaf AC.

Refer to [Appendix B](#) for different data forwarding scenarios.

5. Backward Compatibility

5.1. AC Type

Since there is no restriction on communication between any ACs connected to a VPLS instance as per current standard, on a PE not supporting the extension specified in this document, an AC is functionally equivalent to a Root AC as per the extension specified in this document.

5.2. Control Word

For backward compatibility reason, use of Control Word on Ethernet PW is not mandatory. If the PE on the other end prefers not to use Control Word or does not support Control Word (which implies such PE does not support the extension specified in this document), then Control Word will not be used on the Ethernet PW.

Refer to [Appendix A](#) for different Control Word scenarios.

5.3. Additional Action in Data Forwarding

For backward compatibility reason, a PE not supporting the extension specified in this document can participate in an E-Tree construct, but Leaf ACs MUST NOT be connected to such PE.

Lack of Control Word L-bit per-payload signaling between a PE supporting the extension specified in this document and a PE not supporting such extension does not result in any problem.

Refer to [Appendix B](#) for different data forwarding scenarios.

5.3.1 Ingress PE Support the Extension but Egress PE Not

If Control Word is used on the Ethernet PW, the ingress PE will set CW L-bit to either 0 or 1 (Rule 3 and Rule 4 in Table 1). The egress PE will ignore the reserved bits (which include the L-bit position) as per current standard [[RFC4448](#)] and forward the frame. This is correct as only Root ACs exist on a PE not supporting the extension specified in this document. A Root AC can receive from any other ACs.

If Control Word is not used on the Ethernet PW, there will be no CW L-bit. The egress PE will forward the frame as per current standard. This is correct as only Root ACs exist on a PE not supporting the extension specified in this document. A Root AC can receive from any other ACs.

5.3.2 Egress PE Support the Extension but Ingress PE Not

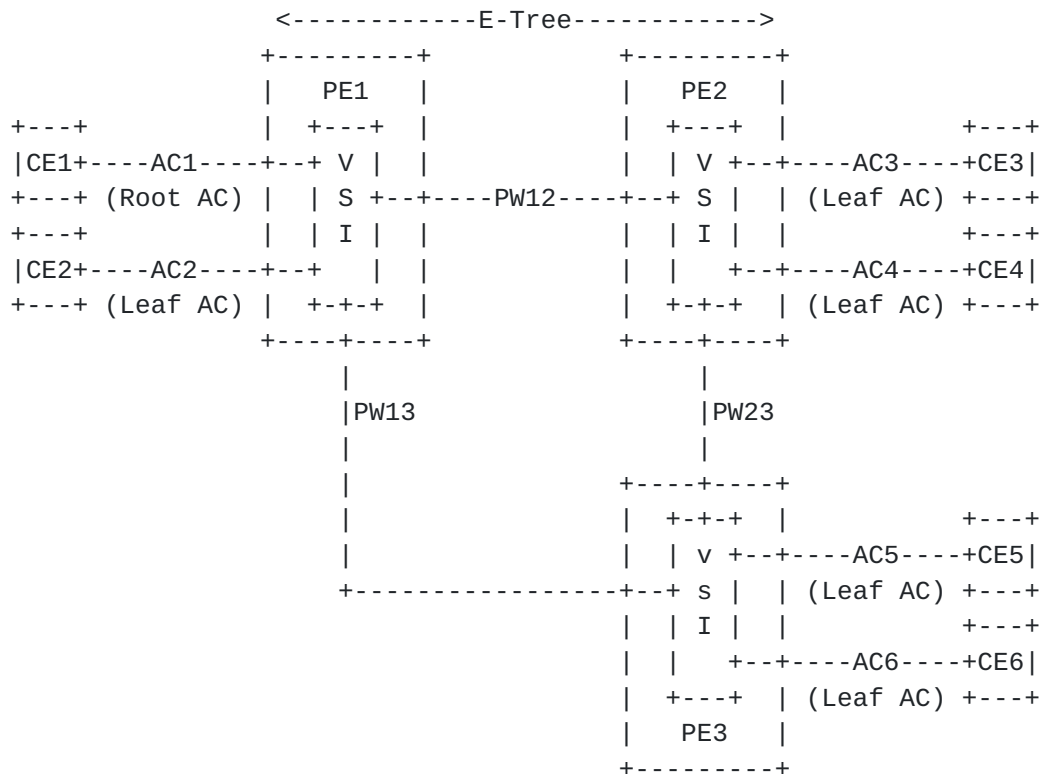
If Control Word is used on the Ethernet PW, the ingress PE will set the reserved bits (which include the L-bit position) to 0 as per current standard [RFC4448]. The egress PE will process it as CW L-bit = 0, which means the frame is from a Root AC. This is correct as only Root ACs exist on a PE not supporting the extension specified in this document. A Root AC can transmit to any other ACs.

If Control Word is not used on the Ethernet PW, there will be no CW L-bit. The egress PE will process it in the same way as CW L-bit = 0 (compare Rule 7 and Rule 8 in Table 1), which means the frame is from a Root AC. This is correct as only Root ACs exist on a PE not supporting the extension specified in this document. A Root AC can transmit to any other ACs.

6. Optional Enhancements for Leaf-only PE

It is important to note that the enhancements specified in this section are OPTIONAL for achieving an E-Tree implementation using VPLS.

In the context of a specific VPLS instance, a Leaf-only PE means that the PE only has Leaf ACs connected to it. In Figure 2 below, PE2 and PE3 are Leaf-only PEs.



<-----E-Tree----->

Figure 2: Reference Model for Leaf-only PE

6.1. No PW between Leaf-only PEs

With an E-Tree construct, a Leaf AC cannot receive from and transmit to any other Leaf ACs. Pseudowire PW23 is between two Leaf-only PEs, and is therefore not supposed to carry any traffic. If removal of such pseudowire can bring significant benefits, enhancement in this aspect may be desirable.

6.2. Not Forward Frame from Leaf AC to Leaf-only PE

When PE1 receives a Broadcast frame via AC2, PE1 will set CW L-bit to 1 and forward the frame to pseudowires PW12 and PW13. PE2 receives the frame via PW12 but does not forward it to any ACs. A similar situation occurs for PE3. Network bandwidth is consumed and then the egress PE decides to drop the frame. This applies to any frames (Broadcast, Multicast, Unicast Known, Unicast Unknown) from a Leaf AC towards a Leaf-only egress PE. If such traffic pattern is significant in volume, enhancement in this aspect may be desirable.

Each PW on a PE MUST have an Leaf-only attribute, either Yes or NO, indicating whether the peer PE at the other end of the PW is a Leaf-only PE or not. The default value for Leaf-only attribute MUST be No.

A PE MUST NOT forward a data frame on a PW with Leaf-only = Yes if any of the following conditions is true

- The frame is from a Leaf AC
- The frame is from PW and the CW L-bit = 1

The Leaf-only attribute for each PW can be locally configured on a PE for each PW, or MAY be decided through signaling between PEs.

For LDP VPLS, Interface Parameter Sub-TLV can be used for such signaling, refer to [section 5.5 in \[RFC4447\]](#).

For BGP VPLS, Signaling PE Capability can be used for such signaling, refer to [section 3.2.4 in \[RFC4761\]](#).

Further details will be added in later version of this document.

7. Hierarchical VPLS

7.1. Hierarchical LDP VPLS

7.1.1. Spoke Connectivity for Bridging-Capable Devices

This refers to [section 10.1.1 in \[RFC4762\]](#).

MTU-s MUST support the extension specified in [section 4](#) of this

document.

Key, et al.

Expires October 2012

[Page 10]

PE-rs MUST support the extension specified in [section 4](#) of this document.

PE-rs MUST perform the additional actions specified in Table 2 below when forwarding a data frame from one PW to another PW.

The "Set CW L-bit" action and "Forward or Drop" decision are in addition to and performed after the following

- MAC-based forwarding decision as per current standard
- Loop free VPLS "split horizon" rule (MUST NOT forward traffic from one PW to another PW in the same VPLS mesh) as per current standard

+-----+-----+-----+-----+-----+-----+-----+						
IF Conditions AND			THEN Actions			
Rule						
	Forward	Receive	CW	Set CW	Forward	
	Frame to PW	Frame from PW	L-bit	L-bit	or Drop	
H1	PW with CW	PW with CW	any	Copy L-bit	Forward	
H2	PW with CW	PW with no CW	n/a	Set to 0	Forward	
H3	PW with no CW	any PW	any	n/a	Forward	

Table 2: Additional Action in Data Forwarding - H-VPLS

[7.1.2.](#) Multi-domain VPLS Service

This refers to [section 10.3 in \[RFC4762\]](#).

Border PE MUST support the extension specified in [section 4](#) of this document.

Border PE MUST perform the additional actions same as those specified for PE-rs in [section 7.1.1](#) when forwarding a data frame from one PW to another PW.

[7.1.3.](#) Spoke Connectivity for Non-Bridging Devices

This refers to [section 10.1.3 in \[RFC4762\]](#).

No change is required for PE-r.

Control Word is not required for the spoke PW between PE-rs and PE-r.

PE-rs MUST treat each spoke PW to PE-r equivalent to an AC of the VPLS instance (consider it as an AC extended by a P2P PW), and MUST support the extension specified in [section 4.1](#) (AC Type) and [section 4.3](#) (Additional Action in Data Forwarding) for such spoke PWs in the same way as ACs.

7.1.4. Hierarchical VPLS Model using Ethernet Access Network

This refers to [section 11 in \[RFC4762\]](#).

This will be added in later version of this document.

7.2. Hierarchical BGP VPLS

This refers to [section 3.6 in \[RFC4761\]](#).

No change is required.

8. Compliance with Requirements

This refers to [Draft VPLS ETree Req] [Section 5](#). Requirements.

The solution prohibits communication between any two Leaf ACs in a VPLS instance.

The solution allows multiple Root ACs in a VPLS instance.

The solution allows Root AC and Leaf AC of a VPLS instance co-exist on any PE.

The solution is applicable to both BGP-VPLS [[RFC4761](#)] and LDP-VPLS [[RFC4762](#)].

The solution is applicable to Case 1: Single technology "VPLS Only".

The solution has minimal impact on existing VPLS solution.

The solution is backward compatible with the existing VPLS solution.

9. Security Considerations

This will be added in later version of this document.

10. IANA Considerations

This document does not require IANA assignment.

11. Acknowledgements

The authors would like to thank Chen Ran, Weilian Jiang and Yuji

Kamite for their valuable comments.

Key, et al.

Expires October 2012

[Page 12]

12. References

12.1. Normative References

- [RFC2119] Bradner, S., Key words for use in RFCs to Indicate Requirement Levels, [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4385] Bryant, S., Swallow, G., and Al, Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN, February 2006.
- [RFC4447] Martini, L., and al, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP), April 2006
- [RFC4448] Martini, L., and al, Encapsulation Methods for Transport of Ethernet over MPLS Networks, April 2006
- [RFC4761] Kompella & Rekhter, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling, January 2007
- [RFC4762] Lasserre & Kompella, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling, January 2007
- [MEF6.1] Metro Ethernet Forum, Ethernet Services Definitions - Phase 2, April 2008

12.2. Informative References

- [Draft VPLS ETree Req]
Key, et al., Requirements for MEF E-Tree Support in VPLS, [draft-ietf-l2vpn-etree-reqt-00](#) (work in progress), October 2011
- [Draft CW L-bit]
Delord, et al., Control Word Reserved bit for use in E-Tree, [draft-delord-pwe3-cw-bit-etree-07](#) (work in progress), April 2012

Appendix A. Control Word Scenarios

As per current standard [[RFC4447](#)]: If both endpoints prefer the use of the control word, this procedure will cause it to be used. If either endpoint prefers not to use the control word or does not support the control word, this procedure will cause it not to be used. If one endpoint prefers to use the control word but the other does not, the one that prefers not to use it has no extra protocol to execute; it just waits for a Label Mapping message that has c=0.

There are three possible scenarios.

Scenario	PE1 Support this Extension?	PE2 Support this Extension?	Support Control Word?	Control Word used on PW between PE1 and PE2?
1	Yes	Yes	Yes	Yes
2	Yes	No	Yes	Yes
3	Yes	No	No	No

Table 2: Control Word Scenarios

For Scenario 2, although Control Word is used on the Ethernet PW, the two PEs process the Control Word L-bit differently:

- PE1 will always set and interpret the CW L-bit as specified in [Draft CW L-bit], i.e. 0 = from Root AC; 1 = from Leaf AC.
- PE2 will always set the CW L-bit to 0 when sending a frame on the PW and ignore the CW L-bit when receiving a frame from the PW. Actually PE2 has no concept of CW L-bit but just treat it as bit 4, one of the reserved bits for future use.

Appendix B. Data Forwarding Scenarios

B.1. Reference Model

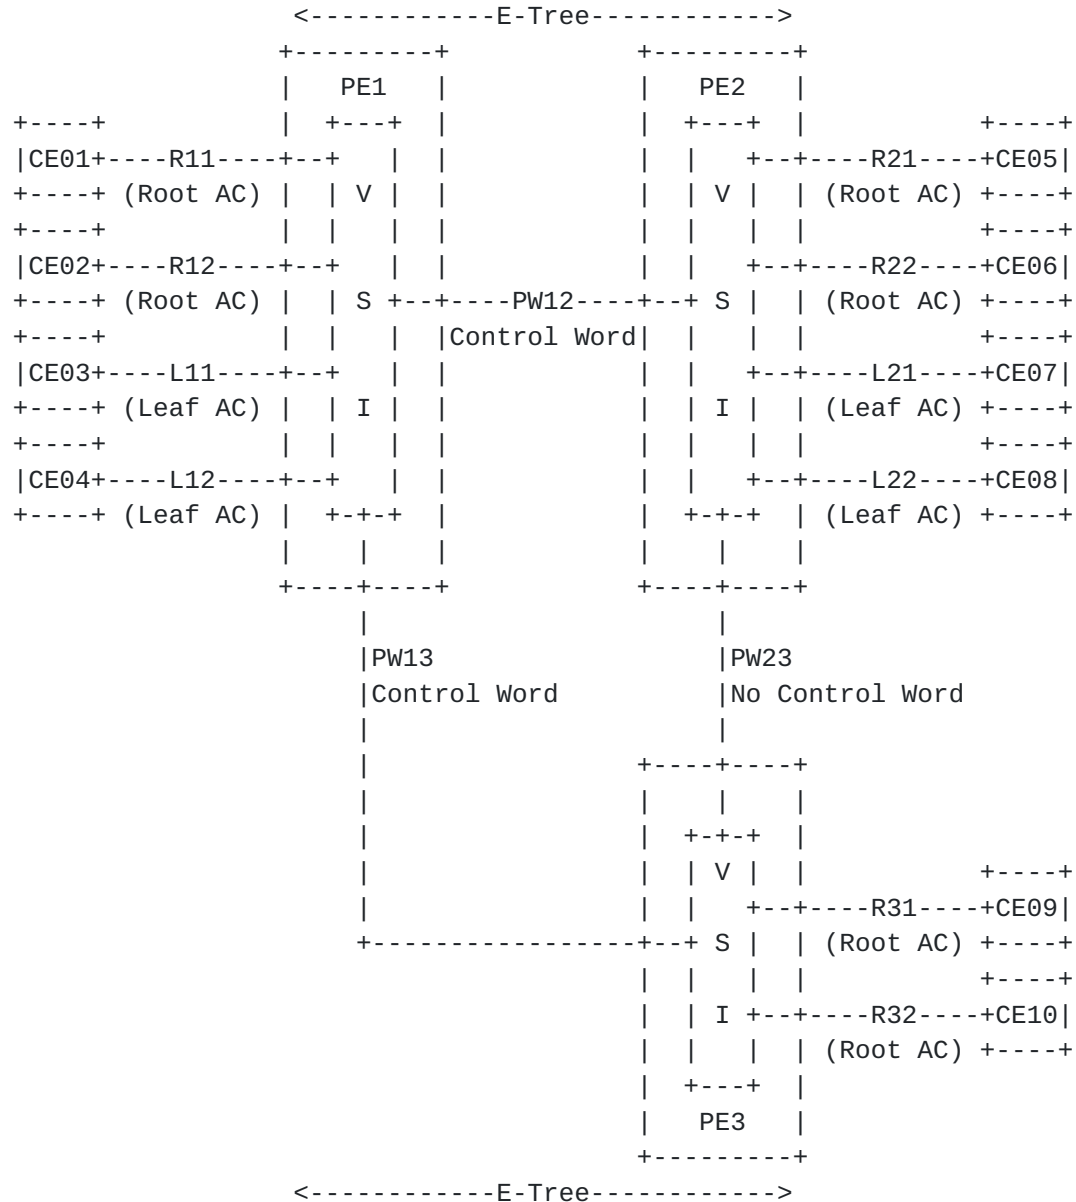


Figure 3: Reference Model for Data Forwarding Scenarios

PE1 and PE2 both support the extension specified in this document.

PE3 does not support the extension specified in this document.

In this appendix

- "this extension" means the extension to VPLS specified in this

document.

- "Rule" refers to Rules in Table 1 in [Section 4.3](#).

Key, et al.

Expires October 2012

[Page 15]

B.2. Unicast Known

B.2.1. From Root AC to Root AC

9 scenarios, per source and destination matching in table below.

Source		Destination		
		PE1	PE2	PE3
PE1	R11	R12	R21	R31
PE2	R21	R11	R22	R31
PE3	R31	R11	R21	R32

B.2.1.1. From R11 to R12 (same PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to R12.
- (2) PE1: Forward as per this extension. (Rule 1)

B.2.1.2. From R11 to R21 (different PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to PW12.
- (2) PE1: As per this extension, set L-bit to 0. (Rule 3)
- (3) PE2: MAC-based forwarding decision as per current standard, forward frame to R21.
- (4) PE2: Forward as per this extension. (Rule 1)

B.2.1.3. From R11 to R31 (different PE, egress PE not support this extension, control word on PW)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to PW13.
- (2) PE1: As per this extension, set L-bit to 0. (Rule 3)
- (3) PE3: MAC-based forwarding decision as per current standard, forward frame to R31.
- (4) PE3: PE3 not support this extension, L-bit ignored and forward frame.

B.2.1.4. From R21 to R11 (different PE)

Similar to B.2.1.2.

B.2.1.5. From R21 to R22 (same PE)

Similar to B.2.1.1.

B.2.1.6. From R21 to R31 (different PE, egress PE not support this extension, no control word on PW)

- (1) PE2: MAC-based forwarding decision as per current standard, forward frame to PW23.
- (2) PE2: No additional L-bit action as per this extension. (Rule 2)
- (3) PE3: MAC-based forwarding decision as per current standard, forward frame to R31.
- (4) PE3: PE3 not support this extension, forward frame.

B.2.1.7. From R31 to R11 (different PE, ingress PE not support this extension, control word on PW)

- (1) PE3: MAC-based forwarding decision as per current standard, forward frame to PW13.
- (2) PE3: PE3 not support this extension, no L-bit action.
- (3) PE1: MAC-based forwarding decision as per current standard, forward frame to R11.
- (4) PE1: Forward as per this extension. (Rule 1)

B.2.1.8. From R31 to R21 (different PE, ingress PE not support this extension, no control word on PW)

- (1) PE3: MAC-based forwarding decision as per current standard, forward frame to PW23.
- (2) PE3: PE3 not support this extension, no L-bit action.
- (3) PE2: MAC-based forwarding decision as per current standard, forward frame to R21.
- (4) PE2: Forward as per this extension. (Rule 1)

B.2.1.9. From R31 to R32 (same PE, PE not support this extension)

- (1) PE3: MAC-based forwarding decision as per current standard, forward frame to R32.
- (2) PE3: PE3 not support this extension, forward frame.

B.2.2. From Root AC to Leaf AC

6 scenarios, per source and destination matching in table below.

Source		Destination		
		PE1	PE2	PE3
PE1 R11	L11	L21	---	
PE2 R21	L11	L21	---	
PE3 R31	L11	L21	---	

B.2.2.1. From R11 to L11 (same PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to L11.
- (2) PE1: Forward as per this extension. (rule 5)

B.2.2.2. From R11 to L21 (different PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to PW12.
- (2) PE1: As per this extension, set L-bit to 0. (Rule 3)
- (3) PE2: MAC-based forwarding decision as per current standard, forward frame to L21.
- (4) PE2: Forward as per this extension. (Rule 8)

B.2.2.3. From R21 to L11 (different PE)

Similar to B.2.2.2.

B.2.2.4. From R21 to L21 (same PE)

Similar to B.2.2.1.

B.2.2.5. From R31 to L11 (different PE, ingress PE not support this extension, control word on PW)

- (1) PE3: MAC-based forwarding decision as per current standard, forward frame to PW13.
- (2) PE3: PE3 not support this extension, no L-bit action. As per current standard, the corresponding bit position must be set to 0.
- (3) PE1: MAC-based forwarding decision as per current standard, forward frame to L11.
- (4) PE1: Forward as per this extension. (Rule 8)

B.2.2.6. From R31 to L21 (different PE, ingress PE not support this extension, no control word on PW)

- (1) PE3: MAC-based forwarding decision as per current standard, forward frame to PW23.
- (2) PE3: PE3 not support this extension, no L-bit action.
- (3) PE2: MAC-based forwarding decision as per current standard, forward frame to L21.
- (4) PE2: Forward as per this extension. (Rule 7)

B.2.3. From Leaf AC to Root AC

6 scenarios, per source and destination matching in table below.

Source		Destination			
		PE1	PE2	PE3	
PE1	L11	R11	R21	R31	
PE2	L21	R11	R21	R31	
PE3	---	---	---	---	

B.2.3.1. From L11 to R11 (same PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to R11.
- (2) PE1: Forward as per this extension. (Rule 1)

B.2.3.2. From L11 to R21 (different PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to PW12.
- (2) PE1: As per this extension, set L-bit to 1. (Rule 4)
- (3) PE2: MAC-based forwarding decision as per current standard, forward frame to R21.
- (4) PE2: Forward as per this extension. (Rule 1)

B.2.3.3. From L11 to R31 (different PE, egress PE not support this extension, control word on PW)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to PW13.
- (2) PE1: As per this extension, set L-bit to 1. (Rule 4)
- (3) PE3: MAC-based forwarding decision as per current standard, forward frame to R31.
- (4) PE3: PE3 not support this extension, as per current standard L-bit ignored and forward frame.

B.2.3.4. From L21 to R11 (different PE)

Similar to B.2.3.2.

B.2.3.5. From L21 to R21 (same PE)

Similar to B.2.3.1.

B.2.3.6. From L21 to R31 (different PE, egress PE not support this extension, no control word on PW)

- (1) PE2: MAC-based forwarding decision as per current standard, forward frame to PW23.
- (2) PE2: No L-bit action as per this extension. (Rule 2)
- (3) PE3: MAC-based forwarding decision as per current standard, forward frame to R31.
- (4) PE3: PE3 not support this extension, forward frame.

B.2.4. From Leaf AC to Leaf AC (MUST NOT Forward)

4 scenarios, per source and destination matching in table below.

Source		Destination			
		PE1	PE2	PE3	
PE1	L11	L12	L21	---	
PE2	L21	L11	L22	---	
PE3	---	---	---	---	

B.2.4.1. From L11 to L12 (same PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to L12.
- (2) PE1: As per this extension, drop frame. (Rule 6)

B.2.4.2. From L11 to L21 (different PE)

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to PW12.
- (2) PE1: As per this extension, set L-bit to 1. (Rule 4)
- (3) PE2: MAC-based forwarding decision as per current standard, forward frame to L21.
- (4) PE2: As per this extension, drop frame. (Rule 9)

B.2.4.3. From L21 to L11 (different PE)

Similar to B.2.4.2.

B.2.4.4. From L21 to L22 (same PE)

Similar to B.2.4.1.

B.3. Unicast Unknown

B.3.1. From Root AC

Forward to all other AC.

B.3.1.1. From R11

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to R12, L11, L12, PW12 and PW13.
- (2) PE1: Forward as per this extension. (Rule 1, 5, 3)
- (3) PE1: For PW12 and PW13, as per this extension, set L-bit to 0. (Rule 3)
- (4) PE2: MAC-based forwarding decision as per current standard, forward frame to R21, R22, L21 and L22.
- (5) PE2: As per this extension, forward frame. (Rule 1, 8)
- (6) PE3: MAC-based forwarding decision as per current standard, forward frame to R31 and R32.
- (7) PE3: PE3 not support this extension, as per current standard L-bit ignored and forward frame.

B.3.1.2. From R21

- (1) PE2: MAC-based forwarding decision as per current standard, forward frame to R22, L21, L22, PW12 and PW23.
- (2) PE2: Forward as per this extension. (Rule 1, 5, 3, 2)
- (3) PE2: For PW12, as per this extension, set L-bit to 0. (Rule 3)
- (4) PE2: For PW23, no L-bit action as per this extension. (Rule 2)
- (5) PE1: MAC-based forwarding decision as per current standard, forward frame to R11, R12, L11 and L12.
- (6) PE1: As per this extension, forward frame. (Rule 1, 8)
- (7) PE3: MAC-based forwarding decision as per current standard, forward frame to R31 and R32.
- (8) PE3: PE3 not support this extension, forward frame.

B.3.1.3. From R31

- (1) PE3: MAC-based forwarding decision as per current standard, forward frame to R32, PW13 and PW23.
- (2) PE3: PE3 not support this extension, forward frame.
- (3) PE3: For PW13, PE3 not support this extension, no L-bit action. As per current standard, the corresponding bit position must be set to 0.
- (4) PE3: For PW23, PE3 not support this extension, no L-bit action. No control word on PW23.
- (5) PE1: MAC-based forwarding decision as per current standard, forward frame to R11, R12, L11 and L12.
- (6) PE1: As per this extension, forward frame. (Rule 1, 8)
- (7) PE2: MAC-based forwarding decision as per current standard, forward frame to R21, R22, L21 and L22.
- (8) PE2: As per this extension, forward frame. (Rule 1, 7)

B.3.2. From Leaf AC

Forward to all Root AC only.

B.3.2.1. From L11

- (1) PE1: MAC-based forwarding decision as per current standard, forward frame to R11, R12, L12, PW12 and PW13.
- (2) PE1: For R11 and R12, forward as per this extension. (Rule 1)
- (3) PE1: For L12, as per this extension, drop frame. (Rule 6)
- (4) PE1: For PW12 and PW13, as per this extension, set L-bit to 1. (Rule 4)
- (5) PE2: MAC-based forwarding decision as per current standard, forward frame to R21, R22, L21 and L22.
- (6) PE2: For R21 and R22, forward as per this extension. (Rule 1)
- (7) PE2: For L21 and L22, as per this extension, drop frame. (Rule 9)
- (8) PE3: MAC-based forwarding decision as per current standard, forward frame to R31 and R32.
- (9) PE3: PE3 not support this extension, as per current standard L-bit ignored and forward frame.

B.3.2.2. From L21

- (1) PE2: MAC-based forwarding decision as per current standard, forward frame to R21, R22, L22, PW12 and PW23.
- (2) PE2: For R21 and R22, forward as per this extension. (Rule 1)
- (3) PE2: For L22, as per this extension, drop frame. (Rule 6)
- (4) PE2: For PW12, as per this extension, set L-bit to 1. (Rule 4)
- (5) PE2: For PW23, no L-bit action as per this extension. (Rule 2)
- (6) PE1: MAC-based forwarding decision as per current standard, forward frame to R11, R12, L11 and L12.
- (7) PE1: For R11 and R12, forward as per this extension. (Rule 1)
- (8) PE1: For L11 and L12, as per this extension, drop frame. (Rule 9)
- (9) PE3: MAC-based forwarding decision as per current standard, forward frame to R31 and R32.
- (10) PE3: PE3 not support this extension, forward frame.

B.4. Broadcast

Same as Unicast Unknown in B.3.

B.5. Multicast

Same as Unicast Unknown in B.3.

Authors' Addresses

Raymond Key
Huawei
Email: raymond.key@ieee.org

Simon Delord
Alcatel-Lucent
Email: simon.delord@gmail.com

Frederic Jounay
Orange CH
4 rue caudray 1020 Renens
Switzerland
Email: frederic.jounay@orange.ch

Wim Henderickx
Alcatel-Lucent
Copernicuslaan 50
2018 Antwerp, Belgium
Email: wim.henderickx@alcatel-lucent.com

Lucy Yong
Huawei USA
1700 Alma Dr. Suite 500
Plano, TX 75075, USA
Email: lucy.yong@huawei.com

Lizhong Jin
ZTE Corporation
889, Bibo Road
Shanghai, 201203, China
Email: lizhong.jin@zte.com.cn

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

