

Network Working Group
Internet Draft
Category: Standard Track
Expires: April 2011

Rafi Ram, Orckit-Corrigent
Amir Halperin, Orckit-Corrigent
Raymond Key, Telstra

October 4, 2010

Extension to LDP-VPLS for E-Tree Using Two PW
draft-ram-12vpn-ldp-vpls-etree-2pw-00.txt

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/lid-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 April 4, 2011.

Abstract

This document proposes a solution for Metro Ethernet Forum (MEF) Ethernet Tree (E-Tree) support in Virtual Private LAN Service using LDP Signaling (LDP-VPLS) [RFC4762]. The proposed solution is characterized by the use of two PWs between a pair of PEs. This solution is applicable for both VPLS and H-VPLS.

Table of Contents

1. Introduction.....	3
2. The Problem.....	3
3. The 2-PW Solution.....	3
4. Extension to VPLS for E-Tree.....	4
4.1. AC E-Tree Type.....	4
4.2. PW VSI Type.....	4
4.3. Additional Filtering in Data Forwarding.....	4
4.4. Root/Leaf PWs Signaling.....	5
4.5. Supporting Remote ACs.....	5
5. Backward Compatibility.....	6
6. Compliance with Requirements.....	6
7. Security Considerations.....	6
8. IANA Considerations.....	6
9. Acknowledgements.....	6
10. References.....	7
10.1. Normative References.....	7
10.2. Informative References.....	7
Authors' Addresses.....	7
Copyright Notice.....	7

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

This document proposes a solution for Metro Ethernet Forum (MEF) Ethernet Tree (E-Tree) support in Virtual Private LAN Service using LDP Signaling (LDP-VPLS) [RFC4762].

[Draft ETree VPLS Req] is used as requirement specification.

The proposed solution is characterized by the use of two PWs between a pair of PEs, which requires extension to the current VPLS standard [RFC4762].

This solution is applicable for both VPLS and H-VPLS.

The proposed solution is composed of three main components:

- Current standard LDP-VPLS [RFC4762]
- Extension to LDP-VPLS specified in this document
- PE local split horizon mechanism

2. The Problem

[Draft ETree VPLS Req] identifies the problem when there are two or more PEs with both Root AC and Leaf AC.




Figure 1: Problem Scenario for Leaf-to-Leaf Communication Restriction

- When PE2 receives a frame from PE1 via the Ethernet PW,
- PE2 does not know whether the ingress AC is a Leaf AC or not
 - PE2 does not have sufficient information to enforce the Leaf-to-Leaf communication restriction

3. The 2-PW Solution

A simple fix is to carry additional information with each frame on the PW, indicating whether the frame is originated from a Leaf AC or a Root AC on the ingress PE.

This solution uses a pair of PWs between a pair of VPLS PEs.

- First PW is used for frames originated from Root ACs
- Second PW is used for frames originated from Leaf ACs




Figure 2: Two-PW Solution for Leaf-to-Leaf Communication Restriction
 Extension to current VPLS standard [RFC4762] is required.

4. Extension to VPLS for E-Tree

4.1. AC E-Tree Type

Each AC connected to a specific VPLS instance on a PE MUST have an AC E-Tree Type attribute, either Leaf AC or Root AC. For backward compatibility, the default AC E-Tree Type MUST be Root.

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

4.2. PW VSI Type

Use a pair of PWs to interconnect between a pair of VPLS PEs:
 - Root PW use VSI Root PW type
 - Leaf PW use VSI Leaf PW type

Both PWs SHALL use the same PW-id. This is for indicating in the local node that a pair of such PWs are part of the same logical VSI interface.

On reception, the two PWs SHALL be handled as the same logical VSI interface with respect to MAC address learning/forwarding, e.g. traffic SHALL not be forwarded between such PWs, MAC addresses arriving from one of the PWs SHALL be learnt with a common logical VSI interface.

The VPLS processing entity SHALL send Root originated traffic via the Root PW, and SHALL send Leaf originated traffic via the Leaf PW.

4.3. Additional Filtering in Data Forwarding

An egress PE SHALL NOT deliver a frame originated from a Leaf AC to another Leaf AC.

The following specifies how AC E-Tree type per frame is determined:

- A frame received from a root PW indicates that the frame was originated from a Root AC.
- A frame received from a leaf PW indicates that the frame was originated from a Leaf AC.
- For the case where both ingress AC and egress AC are on the same PE, local split horizon implementation on the PE will be sufficient, and is not further discussed in this document.

4.4. Root/Leaf PWs Signaling

Signaling of root and leaf PWs is require only when two PWs are used for interconnecting between pair of VSIs. For this purpose a new PW types SHALL be used:

- VSI Root PW type SHALL be used to signal a root PW.
- VSI Leaf PW type SHALL be used to signal a leaf PW.

PW type signaling rules remains as defined in [RFC4447].

4.5. Supporting Remote AC

When PW is used to interconnect between VSI and a remote AC (e.g. the PW1, PW2 in Figure 3), a Ethernet Raw or Ethernet tagged PW types SHALL be used as defined in [RFC4762].

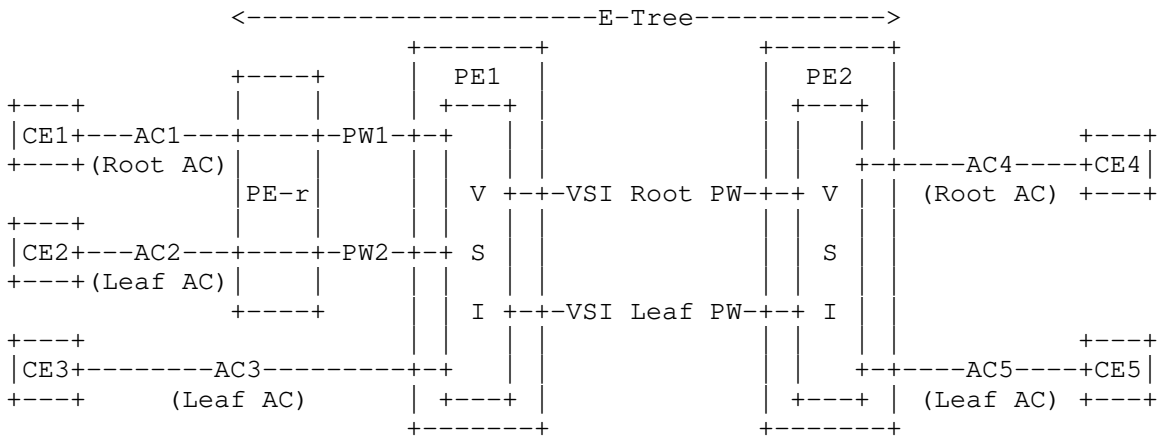


Figure 3: VPLS with Remote AC Connectivity

In addition, the AC type i.e. Root or leaf, SHALL be locally provisioned on the VSI side to specifies the remote AC E-Tree Type per PW. Moreover, such PWs that are used for interconnecting between a remote AC and a VSI SHALL considered as separate logical VSI interfaces with respect to MAC address learning/forwarding e.g. traffic forwarding between such PWs is allowed as long as they are not both defined as Leaf.

In Figure 3, AC1 is remotely interconnected to the VPLS service via PW1, and AC2 is remotely interconnected to the VPLS service via PW2.

AC1 is a Root AC and therefore the local type for PW1 in PE1 SHALL be Root.

AC2 is a Leaf AC and therefore the local type for PW2 in PE1 SHALL be Leaf.

5. Backward Compatibility

VSI root or VSI leaf PW type SHALL be used only in cases where both PEs are VPLS capable and both supports E-Tree root/leaf.

In a case where one of the peers do not support E-Tree, an Ethernet Raw or Ethernet tagged PW types SHALL be used as defined in [RFC4762].

6. Compliance with Requirements

This refers to [Draft ETree VPLS 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 LDP-VPLS [RFC4762].

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

7. Security Considerations

This will be added in later version.

8. IANA Considerations

Additional assignments will be required for the new MPLS PW types introduced in Section 4.2. Details will be added in later version.

9. Acknowledgements

This will be added in later version.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., Key words for use in RFCs to Indicate Requirement Levels, BCP 14, RFC 2119, March 1997.
- [RFC4447] Martini, L., and al, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP), April 2006
- [RFC4762] Lasserre & Kompella, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling, January 2007

10.2. Informative References

- [Draft VPLS ETree Req]
Key, et al., Requirements for MEF E-Tree Support in VPLS, draft-key-l2vpn-vpls-etree-reqt-01.txt, September 2010

Authors' Addresses

Rafi Ram
Orckit-Corrigent
126 Yigal Alon st.
Tel Aviv, Israel
Email: rafir@orckit.com

Amir Halperin
Orckit-Corrigent
126 Yigal Alon st.
Tel Aviv, Israel
Email: amirh@orckit.com

Raymond Key
Telstra
242 Exhibition Street, Melbourne
VIC 3000, Australia
Email: raymond.key@team.telstra.com

Copyright Notice

Copyright (c) 2010 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.

Operations and Management Area Working Group
Internet Draft
Intended status: Standard Track
Expires: Sept 2011

N. So
Verizon
P. Unbehagen
Alcatel-Lu
L. Dunbar
Huawei
H.Yu
TW Telecom
J. Heinz
CenturyLink
N.Figueira
Brocade
March 7, 2011

Requirement and Framework for VPN-Oriented Cloud Services

draft-so-vpn-o-cs-00.txt

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/lid-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 September 7, 2011.

Copyright Notice

Copyright (c) 2011 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the BSD License.

Abstract

This contribution addresses the service providers' requirements to support VPN-Oriented Cloud services. It describes the characteristics of VPN-oriented Cloud Service and specifies the requirement on how to maintain and manage the data center resources for those services.

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 RFC-2119 Error! Reference source not found..

Table of Contents

1. Introduction	3
2. Terminology	3
3. Service definitions and requirements	4
4. Requirements of Data Center networks in support of VPN-Oriented Cloud Services	5
5. Data Center Resource Management Requirements for VPN-oriented Cloud Service	6
6. Security Requirement	7
7. Other Requirements	7
8. IANA Considerations	8
9. Acknowledgments	8
10. References	8
Authors' Addresses	8
Intellectual Property Statement.....	9
Disclaimer of Validity	9

1. Introduction

Layer 2 and 3 VPN services offer secure and logically dedicated connectivity among multiple sites for enterprises. VPN-oriented Cloud Service is for those VPN customers who want to offload some dedicated user data center operations such as software, compute, and storage, to the shared cloud centers. Those customers often do not feel comfortable using public Internet as the cloud center access network. They also have more restrictive requirements on what and how the virtualized cloud center resources, e.g., computing power, disk spaces, and/or application licenses, can be shared.

VPN-Oriented Cloud Services allow the VPN services to be extended into cloud data centers and to control the virtual resources sharing functions. As a network and cloud service provider, a VPN-Oriented Cloud-service product may be offered globally across multiple data centers. Some of the data centers may be owned by a network provider, while others may be owned by a partner/vendor. In addition, multiple VPN-oriented Cloud-Service products can be offered from the same data center.

VPN-Oriented Cloud Services differentiate itself from other cloud services in the following aspects:

- Strictly maintaining the secure, reliable, and logical isolation characteristics of VPN;

- Making the traditional data center services (like computing, storage space, or application licenses) as additional attributes to VPNs.

- VPN having the control on how and what data center resources to be associated with the VPN.

This draft describes the characteristics of those services, their service requirements, and the corresponding requirements to data center networks. It also describes a list of the problems that this service is causing to the network provider/operator, especially for the existing VPN customers. These issues must be addressed immediately in order for service providers to facilitate the addition of Cloud-based services to the VPNs of existing customer.

2. Terminology

DC: Data Center

VM: Virtual Machines

VPN: Virtual Private Network

3. Service definitions and requirements

There are various types of VPN-Oriented Cloud Services. Here are just some examples:

VPN-oriented cloud computing service

This refers to Virtual Machines (VMs) and/or physical servers in a cloud data center being added to a VPN customer. The VPN customer can choose different properties on the computing power, such as dedicated servers, preference on which data center to host those servers, or special VMs which are shared with a group of other VPN customers, and etc.

Any cloud data center providing the VPN-oriented computing services SHOULD be able to automatically provision and/or change the required resources based on the specified properties associated with a VPN.

VPN customers SHALL be able to automatically instantiate or remove hosts to/from the VPN's associated Virtual Machines or dedicated servers through the changing of the customer's VPN properties.

VPN-oriented cloud storage service

This refers to disk space, either virtual or actual blocks of hard drives in data centers, being added to a customer's VPN. The VPN customer SHOULD be able to choose different properties on the storage space, such as: if the content has to be replicated locally or has to be replicated at geographically different locations; if the storage has to be co-located with certain hosts; or which hosts have access to the content, and etc.

These properties are strictly associated with the VPN. Any data center providing the storage space for a VPN SHOULD be able to automatically provision or change the required storage space based on the property associated with the VPN.

The VPN customer SHOULD be able to automatically add disc space or remove disc space to the VPN's associated storage through the changing of the VPN properties.

Each VPN SHALL have the ability to limit the mobility of the stored data to a certain geographic region confinement (country/state).

4. Requirements of Data Center networks in support of VPN-Oriented Cloud Services

The success of VPN services in the enterprise and the government world is largely due to its ability to virtually segregate the customer traffic at layer 2 and layer 3. The lower the layer that segregation can be maintained, the safer it is for the customers from security and privacy perspectives. Today's Data Centers use VLANs to segregate servers and traffic from different customers. Since each customer usually needs multiple zones (e.g., DMZ, Web Server zone, and etc) to place different applications, each customer usually needs multiple VLANs. Even small data centers today already consume several thousands of VLANs. Therefore, pure VLAN segregation is not enough for large data centers.

Network service providers view data center resources as added attributes to VPNs. Therefore, traffic segregation per VPN is an essential requirement to the success of VPN-oriented Cloud-Services in the enterprises and government markets. Other essential requirements include:

Requirements for extending VPNs into data center networks using VPN gateways:

- o The Cloud Service associated with certain VPN(s) SHALL be transmitted over a pre-defined set of connections, and each VPN utilizing the service SHALL be transmitted over a sub-set of logical connections.
- o The VPN gateway should maintain a mapping among Virtual or physical Resources, physical/logical connections, with specific VPNs.
- o The VPN Gateway SHOULD be able to control the connection traffic flow and assign the dedicated virtual resources accordingly.

Independent of the L2/3 technology, e.g., TRILL, PBB, SPB, OpenFlow, and etc, used for connecting external (customer) VPNs and data center virtual resources, e.g., , each VPN SHALL be given a unique Service ID, and traffic separation SHALL be maintained per Service ID.

When a L2/3 VPN is used as the network technology connecting the external (customer) VPN and the data center virtual resources, each external VPN SHALL be mapped to a unique internal VPN.

5. Data Center Resource Management Requirements for VPN-oriented Cloud Service

Today, data center server resources are managed by data center servers' administrators or management systems, and supported by hypervisors on the servers. The entire process is invisible to the underlying networks. The data center management functions today include managing servers, instantiating hosts to VMs, managing disk space, and etc.

Traffic loading and balancing and QoS assignments for data center networks are usually not considered by Data Center's server administration systems. There shall be a way that the VPN can connect with the Data Center's server administration systems that are important to the concept and spirit of the VPN:

The resources in data center MUST be partitioned per VPN's requirements instead of the traditional partitioning per customer. The Cloud orchestration system SHALL have the ability to dedicate a specific block of disk space per services per VPN.

If a VPN requires dedicated access to blocks of disk space, the data center disk management system SHALL allocate the required disk space per VPN and be able to let VPN automatically retrieve the identification of those disk spaces.

If a VPN specifies its associated storage space to be accessible only by certain hosts, the data center disk management system SHALL have the ability to indicate the mechanism used to prevent the unwanted data retrieval for the block of disk space after it is no longer used by the VPN, before it can be re-used by other parties.

The VPN SHALL have the ability to request dedicated L2/3 network resources within the data center such as bandwidth, priorities, and so on.

The VPN SHALL have the ability to hold the requested resources without sharing with any other parties.

The VPN's QoS assignments SHOULD be able to synchronize with the Cloud virtual resources' QoS assignments.

6. Security Requirements

VPN-Oriented Cloud Service SHOULD support a variety of security measures in securing tenancy of virtual resources such as resource locking, containment, authentication, access control, encryption, integrity measure, and etc.

The VPN-Oriented Cloud Service SHOULD allow the security to be configured end-to-end on a per VPN per-user basis. For example, the Virtual Systems MUST resource-lock resources such as memory, but must also provide a cleaning function to insure confidentiality before being reallocated.

VPN-Oriented Cloud Service for private Clouds SHOULD specify an authentication mechanism based on an authentication algorithm (MD5, HMAC-SHA-1) for both header and payload. Encryption MAY also be used to provide confidentiality.

Security boundaries MAY also be create to maintain domains of TRUSTED, UNTRUSTED, and Hybrid. Within each domain access control, techniques MAY be used to secure resources and administrative domains.

7. Other Requirements

The VPN-Oriented Cloud Service SHALL support automatic end-to-end network configuration.

The VPN-Oriented Cloud Service solution MUST have sufficient OAM mechanisms in place to allow consistent end-to-end management of the solution in existing deployed networks. The solution SHOULD use existing protocols (e.g., IEEE 802.1ag, ITU-T Y.1731, BFD) wherever possible to facilitate interoperability with existing OAM deployments.

8. IANA Considerations

9. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

10. References

Authors' Addresses

Ning So
Verizon Inc.
2400 N. Glenville Ave.,
Richardson, TX75082
ning.so@verizonbusiness.com

Paul Unbehagen
Alcatel-Lucent
8742 Lucent Boulevard
Highlands Ranch, CO 80129
paul.unbehagen@alcatel-lucent.com

Linda Dunbar
Huawei Technologies
1700 Alma Drive, Suite 500
Plano, TX 75075, USA
Linda.dunbar@huawei.com

Henry Yu
TW Telecom
10475 Park Meadows Dr.
Littleton, CO 80124
Henry.yu@twtelecom.com

John M. Heinz
CenturyLink
600 New Century PKWY
KSNCAA0420-4B116
New Century, KS 66031
john.m.heinz@centurylink.com

Norival Figueira
Brocade Networks

130 Holger Way
San Jose, CA 95134
nfigureir@brocade.com

Intellectual Property Statement

The IETF Trust takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in any IETF Document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights.

Copies of Intellectual Property disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

All IETF Documents and the information contained therein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

