Network Working Group Internet Draft Category: Standard Track Expires: November 2010

Yaakov(J) Stein RAD Data Communications

Cao Wei Huawei Simon Delord (Ed.) Telstra

> Frederic Jounay Philippe Niger France Telecom

Matthew Bocci Mustapha Aissaoui Alcatel-Lucent

Luca Martini Cisco Systems Inc.

May 24, 2010

# LDP extensions for AII reachability draft-ietf-pwe3-ldp-aii-reachability-04.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

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 <a href="http://www.ietf.org/ietf/lid-abstracts.txt">http://www.ietf.org/ietf/lid-abstracts.txt</a>.

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

This Internet-Draft will expire on November 25, 2010.

Delord et al.

Expires November 2010

[Page 1]

## Abstract

The dynamic End-to-End Multisegment pseudowire setup requires PEs to maintain a pseudowire routing table when using FEC129. There is a requirement to automatically advertise Attachment Individual Identifiers to enable the pseudowire routing tables to be populated. Two mechanisms already exist, a BGP reachability information distribution mechanism and an IGP based one. Here we define a third solution relying on LDP. It allows for automatic advertisement of the Attachment Individual Identifier prefixes provisioned on a T-PE when this node does not run BGP or IGP. The mechanism described here runs on the T-LDP (Targeted LDP) session between the T-PE and S-PE, and is intended to complement existing PW routing mechanisms using BGP or OSPF.

# 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 [<u>RFC2119</u>].

# Table of Contents

<u>1</u> . Introduction <u>3</u>
2. Scope and Applicability3
<u>2.1</u> . Scope <u>3</u>
<u>2.2</u> . Applicability <u>4</u>
3. LDP Extensions5
3.1. LDP AII Address Family Type5
3.2. LDP AII Reachability Capability Advertisement <u>6</u>
4. AII Reachability Advertisement Procedures <u>6</u>
<u>4.1</u> . Detailed AII Address Message Procedures <u>6</u>
<u>4.1.1</u> . T-PE Procedures <u>7</u>
<u>4.1.2</u> . S-PE Procedures <u>7</u>
5. Security Considerations
<u>6</u> . IANA Considerations <u>7</u>
<u>6.1</u> . Address Family Type <u>7</u>
<u>6.2</u> . TLV TYPE NAME SPACE <u>8</u>
<u>7</u> . Acknowledgments
<u>8</u> . References
<u>8.1</u> . Normative References <u>8</u>
<u>8.2</u> . Informative References <u>8</u>
Authors' Addresses <u>8</u>
Intellectual Property and Copyright Statements

Delord et al. Expires November 2010

[Page 2]

## **<u>1</u>**. Introduction

This document describes procedures for PEs to populate their Pseudowire (PW) routing tables via the Label Distribution Protocol (LDP). It is targeted at MultiSegment Pseudowire (MS-PW) applications. The dynamic End-to-End MS-PW architecture requires T-PEs and S-PEs to maintain a PW routing tables when using FEC129. The procedure addresses MS-PW architectures where a T-PE does not (for whatever reason) run either an Interior Gateway Protocol (IGP) or Multi-Protocol Border Gateway Protocol (MP-BGP). The mechanism described here is intended to complement other existing PW routing mechanisms already described in [DYN MS-PW], [BGP AD] and [OSPF MS-PW].

The need for MS-PWs is presented in [<u>RFC5254</u>]. To allow for minimal manual intervention/configuration, FEC129 needs to be used as per [<u>MS-PW</u>] and [DYN MS-PW]. [<u>RFC5003</u>] describes the AII type and the fields that identify pseudowire endpoints called attachment individual identifiers (AII).

[DYN MS-PW] specifies a mechanism, based on the MP-BGP, that enables the advertisement of MS-PW endpoint information in the form of aggregated AIIs through the network, and thus allows the automatic placement of MS-PWs.

[OSPF MS-PW] is another way of enabling the automatic placement of MS-PWs across an OSPF domain when MP-BGP is not / cannot be used (e.g. when PWE3 is extended into the access part of the network).

## **2**. Scope and Applicability

### <u>2.1</u>. Scope

One important application is the use of this LDP protocol extension in access networks that use IP/MPLS as their access technology and MS-PW to support L2 services (as per [<u>RFC5254</u>]). MP-BGP or an IGP is often not available in this part of the network.

< PW Se	gt 1>< PW Seg	t 2>< PW S	egt 3>
++	++	++	++
T-PE1+	+S-PE1+	+S-PE2+	+T-PE2
++	++	++	++
<-static-IP-r	outing> <igf< td=""><td>/MP-BGP-available</td><td>&gt;</td></igf<>	/MP-BGP-available	>

Figure 1 MS-PW Use Case

Figure 1 describes a typical use case where T-PE1 and T-PE2 need to establish one or several MS-PWs between them. A MS-PW will be

Delord et al. Expires November 2010 [Page 3]

composed of at least two PW segments (PW Segt 1 between T-PE1 and S-PE1, PW Segt 2 between S-PE1 and S-PE2 and PW Segt 3 between S-PE2 and T-PE2). However T-PE1 is not running either an IGP or MP-BGP and only has one static default route and a T-LDP session towards SPE1. SPE1, SPE2 and TPE2 are running an IGP and/or MP-BGP.

The aim here is for T-PE1 to announce to the S-PE(s) with which it has a T-LDP session (there may be more than one S-PE) its locally attached PW routes, so that the S-PE(s) can populate their AII PW routing table with the T-PE prefixes.

The AII prefixes locally defined on the T-PE are then advertised via an LDP Address Message containing a new Address Family. This new address family capability will follow the processes defined in [RFC 5561].

This document does not presuppose any specific constraint on the AII PW addressing space (i.e it does not require the AII PW addressing space to be a subset of the IP addressing space). Therefore, this document does not define a routing protocol as such, but rather a "reachability" information distribution method by which a T-PE advertises its AII to a S-PE. The S-PE will then aggregate and advertise this information, using one of the MS-PW dynamic placement mechanisms e.g. MP-BGP, to the other S-PEs and T-PEs in the network. Note that the S-PE will also advertise it's locally attached prefixes, and possibly an optional "default PW route".

### 2.2. Applicability

<u>Section 7.1</u> of [DYN MS-PW] describes 7.1 the different rules for aggregated AII PW routing table lookup. The next signaling hop for a segment of the pseudowire may be determined via a fixed route (static route and typically a "default route").

In the case of MPLS enabled access networks, an S-PE (e.g. a DSLAM or other access node) will aggregate up to a few thousands devices that may require several pseudowires (or "services") on each of those devices.

These devices may not require either an IGP or MP-BGP to be integrated into the network, for example because it may not be desirable for a Service Provider to have to re-engineer their metro/access architecture by extending their IGP or inserting MP-BGP further down in the access network. However, they may already use LDP functionality to setup and maintain pseudowires. They can also be configured with one or two default static routes as described in [DYN MS-PW], however on the S-PE side, the provisioning required becomes increasingly complex. Furthermore, closer to the end node, it is quite possible that some pseudowires will need to be setup, removed or modified (e.g. for a bandwidth upgrade) over a short timescale. Therefore, there is a need for a mechanism that will alleviate the provisioning burden on the S-PE(s).

Delord et al. Expires November 2010 [Page 4]

#### **<u>3</u>**. LDP Extensions

### 3.1. LDP AII Address Family Type

In the case described in this document, we assume LDP sessions between the T-PE and related S-PEs.

A new Address Family (AF) type called "AII address family" (TBA) is defined for the Address-List TLV carried in LDP Address and Withdraw Address messages.

This new AF allows LDP peers to advertise directly attached AII prefixes, as part of an LDP Address Message and to withdraw directly attached AII prefixes as part of an LDP Withdraw Address Message.

When a T-PE needs to advertise a new AII prefix to an S-PE, it sends an LDP Address Message containing the AII prefixes to all the S-PEs with which it maintains LDP sessions.

When a T-PE needs to withdraw an AII, it sends an LDP Withdraw Address Message containing the AII prefixes to all S-PEs with which it maintains LDP sessions.

The Address-List TLV is defined in [<u>RFC5036</u>].

AII address prefix is formatted according to AII Type II, defined in [RFC5003] section 3.2 (figure 1).

0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 |0|0| Address List (0x0101) | Length Address Family | Prefix Length | AII Type II Address (32 - 64) 

Prefix Length One octet unsigned integer containing the length in bits of the address prefix that follows. The minimum prefix length for an AII address MUST be of 32 bits. Prefix lengths of 65 to 95 are invalid as the AC ID field cannot be aggregated.

Two AII Address prefixes (PW routes) are said to match only when both the AII Type II as well as the 8 bits prefix length are the same.

Delord et al. Expires November 2010

[Page 5]

## 3.2. LDP AII Reachability Capability Advertisement

In order to use this method of propagating 12 reachability information a PE must first advertise this capability to all LDP peers. This is achieved by using the methods in [RFC5561] and advertising the AII Address Family LDP capability TLV. If an LDP peer supports the dynamic capability advertisement, this can be done by sending a new capability message with the S bit set for the AII Address Family capability TLV when the first AII Prefix is enabled on the PE. If the peer does not support dynamic capability advertisement, then the AII Address Family TLV MUST be included in the LDP initialization procedures in the capability parameter [RFC 5561].

The following TLV is defined to indicate the AII Address Family capability:

### **<u>4</u>**. AII Reachability Advertisement Procedures

[RFC5036] defines the procedures by which LSRs maintain and exchange label information via LDP.

This document does not change any of the standard LDP procedures; it only adds the AII address family type for the Adress-List TLV carried in LDP Address and Withdraw Address messages.

The rules for advertising and withdrawing addresses are as per [RFC 5036].

#### 4.1. Detailed AII Address Message Procedures

In order for a T-PE to begin advertising its locally attached AII prefixes to its S-PEs, the T-PE must know the address(es) of the remote S-PE(s) and already have a T-LDP with each one of those. This information may have been configured on the T-PE, or it may have been learned dynamically via some autodiscovery procedure.

Delord et al. Expires November 2010

[Page 6]

#### 4.1.1. T-PE Procedures

Upon provisioning the T-PE with a prefix of one or more specific AII(s), the T-PE MUST send an Address Message including its ASN and prefix, and optionally an aggregated AII representing its locally attached AII address(es), to all the S-PEs with which it maintains T-LDP sessions.

The addresses are structured according to AII type II [<u>RFC5003</u>]. The T-PE MUST only advertise its AIIs over its T-LDP sessions towards its directly attached S-PEs.

Whenever an attachment circuit not addressed by an existing aggregated already advertised AII is configured on a T-PE, the T-PE MUST advertise the new address with an Address message.

Whenever a T-PE "de-activates" a previously advertised AII Prefix, it SHOULD withdraw the AII Prefix with an Address Withdraw message.

A T-PE MAY re-advertise an address that it has previously advertised without explicitly withdrawing the address. If the receiver already has this address binding, it SHOULD take no further action.

A T-PE MAY withdraw an address without having previously advertised it. If the receiving S-PE has no address binding, it SHOULD take no further action.

### 4.1.2. S-PE Procedures

If a PE receives an address TLV containing an address family that it does not support, it SHOULD follow the procedures defined in [<u>RFC5036</u>].

An S-PE receiving an address list TLV containing one or more AII addresses should install those addresses in its local PW routing table. It MAY also re-advertise those addresses using any another supported dynamic MS-PW routing protocol like BGP [DYN MS-PW] or OSPF [OSPF MS-PW]. An S-PE MUST NOT send the attached T-PE address to other S-PE using T-LDP, because it is not a local connected address. The only routes that an S-PE MAY redistribute using T-LDP is connected (local) routes and a "default" route as an exception.

#### 5. Security Considerations

[MPLS SEC] provides a security framework for MPLS and GMPLS Networks. It emphasizes LDP as well as inter-provider security considerations. The same security framework and considerations apply to the capability mechanism described in this document.

[Page 7]

# 6. IANA Considerations

#### 6.1. Address Family Type

This document defines a new Address Family type called AII address family (TBA) for the Adress-List TLV carried in LDP Address and Address Withdraw messages.

The following value is suggested for assignment:

Registry Number Description

27 AII Attachment Individual Identifier

### 6.2. TLV TYPE NAME SPACE

This document uses a new LDP TLV type, IANA already maintains a registry of name "TLV TYPE NAME SPACE" defined by [<u>RFC5036</u>]. The following value is suggested for assignment:

TLV Type Description

0x406 AII address family capability TLV

#### 7. Acknowledgments

The authors would like to thank Florin Balus, Wim Hendeickx, Jean-Louis Le Roux, Ed Wong and Raymond Key for their valuable comments and suggestions.

## 8. References

#### 8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, March 1997.
- [RFC5036] Andersson, L., Doolan, P., Feldman, N., Fredette, A., Thomas, B., "LDP Specification", January 2001.
- [RFC5003] Chris Metz et. al., "AII Types for Aggregation", February 2007.

Delord et al. Expires November 2010 [Page 8]

## 8.2. Informative References

- [MS-PW] Martini et al., "Segmented Pseudo Wire", Internet Draft, <u>draft-ietf-pwe3-segmented-pw-14.txt</u>, April 2010
- [DYN MS-PW] Bocci, M., Martini, L., "Dynamic placement of Multi Segment Pseudo Wires", Internet Draft, draft-ietfpwe3-dynamic-ms-pw-10.txt, October 2009
- [BGP AD] E. Rosen et. al., "Provisioning, Autodiscovery, and Signaling in L2VPNs", draft-ietf-l2vpn-signaling-08.txt, May 2006
- [OSPF MS-PW] A. Dolganow, M. Bocci et. al., "OSPF Extensions for Dynamic Multi- segment Pseudo Wires",<u>draft-dolganow-pwe3-ospf-ms-pw-ext-03.txt</u>, February 2008
- [RFC5561] B. Thomas, "LDP Capabilities", July 2009
- [MPLS SEC] L. Fang, "Security Framework for MPLS and GMPLS Networks", draft-ietf-mpls-mpls-and-gmpls-securityframework-09.txt, March 2010
- [RFC5254] Bitar, N., Bocci, M., and Martini, L., "Requirements for inter domain Pseudo-Wires",October 2008

Author's Addresses

Simon Delord Telstra 242 Exhibition St Melbourne VIC 3000 Australia Email: simon.a.delord@team.telstra.com

Frederic Jounay France Telecom 2, avenue Pierre-Marzin 22307 Lannion Cedex FRANCE Email: frederic.jounay@orange-ftgroup.com

Philippe Niger France Telecom 2, avenue Pierre-Marzin 22307 Lannion Cedex

FRANCE Email: philippe.niger@orange-ftgroup.com

Delord et al. Expires November 2010

[Page 9]

May 2010

Mustapha Aissaoui Alcatel-Lucent 600 March Road Kanata ON, Canada e-mail: mustapha.aissaoui@alcatel-lucent.com Matthew Bocci Alcatel-Lucent, Voyager Place Shoppenhangers Road Maidenhead Berks, UK e-mail: matthew.bocci@alcatel-lucent.co.uk Yaakov (Jonathan) Stein RAD Data Communications 24 Raoul Wallenberg St., Bldg C Tel Aviv 69719, Israel EMail: yaakov\_s@rad.com Cao Wei Huawei Technologies Co., Ltd. Huawei Bld., No.3 Xinxi Rd. Shang-Di Information Industry Base Hai-Dian Distinct, Beijing 100085 China Email: caowayne@huawei.com Luca Martini Cisco Systems, Inc. 9155 East Nichols Avenue, Suite 400 Englewood, CO, 80112 e-mail: lmartini@cisco.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. 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 Simplified BSD License.

Delord et al. Expires November 2010 [Page 10]