

Network Working Group
Internet Draft
Category: Standard Track
Expires: March 2009

Simon Delord
Uecomm

Yaakov(J) Stein
RAD Data Communications

Frederic Jounay
Philippe Niger
France Telecom

Cao Wei
Huawei

Matthew Bocci
Mustapha Aissaoui
Alcatel-Lucent

Luca Martini
Cisco Systems Inc.

September 20, 2008

LDP extensions for AII reachability
draft-delord-jounay-pwe3-ldp-aii-reachability-02.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of 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 March 20, 2009.

Internet Draft

LDP AII Reachability

September 2008

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 [[RFC2119](#)].

Table of Contents

1.	Introduction.....	3
2.	Scope and Applicability.....	3
2.1.	Scope.....	3
2.2.	Applicability.....	4
3.	LDP Extensions.....	5
3.1.	LDP AII Address Family Type.....	5
3.2.	LDP AII Reachability Capability Advertisement.....	6
4.	AII Reachability Advertisement Procedures.....	6
4.1.	Detailed AII Address Message Procedures.....	6
4.1.1.	T-PE Procedures.....	7
4.1.2.	S-PE Procedures.....	7
5.	Security Considerations.....	7
6.	IANA Considerations.....	7
6.1.	Address Family Type.....	7
6.2.	TLV TYPE NAME SPACE.....	8
7.	Acknowledgments.....	8

8.	References.....	8
8.1.	Normative References.....	8
8.2.	Informative References.....	8
	Authors' Addresses.....	8
	Intellectual Property and Copyright Statements.....	9

[1.](#) 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 [MS-PW REQ]. To allow for minimal manual intervention/configuration, FEC129 needs to be used as per [MS-PW] and [DYN MS-PW]. [RFC5003] 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

[2.1.](#) 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 [MS-PW REQ]). MP-BGP or an IGP is often not available in this part of the network.




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 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 [LDP CAP].

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

[Section 7.1](#) 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 require basic 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, the closer to the end node, it is quite possible that some pseudowires 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).

[3](#). 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 Address Withdraw 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 Address Withdraw

In order to use this method of propagating l2 reachability information a PE must first advertise this capability to all LDP peers. This is achieved by using the methods in [[LDP-CAP](#)] and advertising the AII Address Family LDP capability TLV. If an LDP peer supports the dynamic capability advertisement, this can be done by

4.1.1. T-PE Procedures

Upon provisioning on the T-PE with a prefix of one or more specific AII(s), the T-PE MUST generate 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 [[RFC5003](#)]. 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 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 [[RFC5036](#)].

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 mechanism.

5. Security Considerations

TBD

6. IANA Considerations

6.1. Address Family Type

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

Address Withdraw messages.

The following value is suggested for assignment:

Delord et al.

Expires March 2009

[Page 7]

Internet Draft

LDP AII Reachability

September 2008

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 [[RFC5036](#)]. 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", [BCP 14](#), [RFC 2119](#), 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.

[8.2.](#) Informative References

- [MS-PW] Martini et al., "Segmented Pseudo Wire", Internet Draft, [draft-ietf-pwe3-segmented-pw-09.txt](#), July 2008

[DYN MS-PW] Bocci, M., Martini, L., "Dynamic placement of Multi Segment Pseudo Wires", Internet Draft, [draft-ietf-pwe3-dynamic-ms-pw-08.txt](#), July 2008

[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", [draft-dolganow-pwe3-ospf-ms-pw-ext-02.txt](#), February 2008

Delord et al.

Expires March 2009

[Page 8]

Internet Draft

LDP AII Reachability

September 2008

[MS-PW REQ] Bitar, N., Bocci, M., and Martini, L., "Requirements for inter domain Pseudo-Wires", Internet Draft, [draft-ietf-pwe3-ms-pw-requirements-07.txt](#), June 2007

[LDP-CAP] B. Thomas, "LDP Capabilities", Internet Draft, [draft-ietf-mpls-ldp-capabilities-02.txt](#), April 2008

Author's Addresses

Simon Delord
Uecomm
8/658 Church St
Richmond VIC 3121
Australia
Email: sdelord@uecomm.com.au

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

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

Delord et al.

Expires March 2009

[Page 9]

Internet Draft

LDP AII Reachability

September 2008

Cao Wei
Huawei Technologies Co., Ltd.
Huawei Bld., No.3 Xinxu Rd.
Shang-Di Information Industry Base
Hai-Dian District, 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

Intellectual Property Statement

The IETF 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 this 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. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR 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 this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein 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 HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Delord et al.

Expires March 2009

[Page 10]

Internet Draft

LDP AII Reachability

September 2008

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

