

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 21, 2012

P. Kwok
P. Dutta
Alcatel-Lucent
F. Jounay
France Telecom
May 20, 2012

Pseudowire Communities
draft-pkwok-pwe3-pw-communities-03

Abstract

[RFC4447] describes a set of procedures for Pseudowire set-up and maintenance using LDP as signaling protocol.

[I-D.ietf-pwe3-dynamic-ms-pw] extends the mechanisms described in [RFC4447] for dynamic placement of multi-segment pseudowires.

This document describes an extension to [RFC4447] procedures which may be used to pass additional information to S-PE/T-PEs when SS-PWs or MS-PWs are set-up.

The intention of the proposed technique is to aid in policy administration, specifically during MS-PW set-up across various S-PEs. The proposed method is very generic so that it can support the management of various parameters or rules while setting up pseudowires with minimal overhead.

Requirements Language

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].

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 21, 2012.

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. 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.

Table of Contents

- [1. Introduction](#) [4](#)
- [2. PW Communities](#) [5](#)
- [3. Defined PW Community Types](#) [6](#)
 - [3.1. PW Template Community](#) [6](#)
 - [3.1.1. PW Generic Template Community](#) [7](#)
 - [3.2. PW Color Community](#) [7](#)
 - [3.2.1. PW Generic Color Community](#) [7](#)
- [4. IANA Considerations](#) [7](#)
- [5. Security Considerations](#) [8](#)
- [6. Acknowledgements](#) [8](#)
- [7. References](#) [8](#)
 - [7.1. Normative References](#) [8](#)
 - [7.2. References](#) [8](#)
- [Authors' Addresses](#) [8](#)

1. Introduction

A Multi-Segment PW (MS-PW) is defined as a set of two or more contiguous segments that behave and function as a single point-to-point PW. An MS-PW enables service providers to extend the reach of PWs across multiple PSN domains.

To facilitate and simplify the control of dynamic MS-PW set-up across S-PEs, this document proposes a grouping or "community" of PWs so that PW set-up decision can also be based on the identity of the group. Such a scheme is expected to significantly simplify dynamic MS-PW signaling [[I-D.ietf-pwe3-dynamic-ms-pw](#)] that controls the MS-PW set-up across the Switching Provider Edge (S-PE) devices.

MS-PW spans across multiple autonomous systems or administrative domains. For security reasons, strict access control is required at S-PEs through which a PW enters another administrative domain. One way is for operators to define a policy at the S-PE that would match the PW set-up requests based on Target Attachment Individual Identifier (TAII) or Source Attachment Individual Identifier (SAII) or Attachment Group Identifier (AGI) etc. Such policies can be complex or very large, leading to administrative overheads or configuration mistakes. Rather, operators could define several tags/colors which can be associated with individual PWs when they are signaled. S-PEs can then apply PW policies based on the received tags, accordingly. This example application eliminates the primary motivation for a complex policy database that may result in the generation of very large PW prefix-based filter rules. A smaller policy database such as this also requires less maintenance, so shortening or eliminating out-of-band maintenance delays.

Another application of PW policies is in underlying transport applications. Each S-PE independently chooses a unidirectional PSN tunnel to map a set of PW segments to their next S-PE or T-PE. Such PSN Tunnels could be Label Distribution Protocol (LDP) [[RFC5036](#)] or Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) [[RFC3209](#)] or Labeled BGP [[RFC3107](#)] based LSPs. There is currently no signaling support in [[I-D.ietf-pwe3-dynamic-ms-pw](#)] to signal a preference for the type of PSN tunnel to bind a PW to at the S-PEs when multiple tunnel types are available. For example, LDP can be preferred over BGP tunnels when both forms of tunnels are available at an S-PE. Secondly, it is also possible that only a specific RSVP tunnel or class of RSVP tunnels based in Admin Groups is preferable to provide a traffic class or QoS treatment, or protection capability, and some form of control is required that LSPs are correctly used by S-PEs. One possible way is to manually configure filter rules by PW ID or AGI/SAII/TAII, but such rules can create significant maintenance overhead and be prone to configuration errors. Further, signaling

each of the various types of PSN tunnel selection criteria/ preferences in PW set-up messages adds significant burden to LDP label mapping procedures.


In Dynamic MS-PW, a T-PE or S-PE may need to choose one next-hop from several Equal Cost Multi-Path (ECMP) next-hops provided by best matching PW Route. One way to do ECMP selection is to apply some form of hash function on AGI/SAII/TAII of the PW but that strictly limits the MS-PW addressing schemes in order to get proper load distribution of MS-PWs across all next-hops. Operators need a predictable way for load balancing MS-PW across ECMP next-hops which is independent of MS-PW addressing schemes.

To address such policy management issues, this draft proposes a very simple solution that allows minimal manual intervention and configuration with no overhead in PW signaling. It introduces a concept of "PW Communities" that can be thought of as templates provisioned at a S-PE/T-PE, based on which of a certain set of rules are applied to all PWs that are tagged as belonging to same community.

Note that PW Community is different from PW Grouping (as defined using PW Group ID) defined in [RFC4447]]. PW Grouping is associated with binding of a set of PWs to a common event group for reduced signaling of various intensive events such as Label withdraw or PW Status Notification etc. However, PW Communities can be thought of a grouping of PWs from policy management perspective. It is not necessary that PW Grouping and PW Communities associated with a PW be correlated.


2. PW Communities

The PW Communities is an OPTIONAL TLV defined as follows which is in the format of LDP TLV [RFC5036].



U/F bits MUST be set to 1. Length is variable. Value field of the TLV contains a set of "PW Communities".

A PW Community is defined as follows:



Type field indicates the specific PW Community Type. The types are introduced to provide a broad classification of various PW communities based on the scope of applicability. Each community type further provides the flexibility to define sub-types within it. Length of a PW community is variable and to be defined by Type and Sub-Type associated with a PW community.

3. Defined PW Community Types

This section introduces a few PW community types and defines the format of the PW Community for those types.

3.1. PW Template Community


A PW Template community (PW Community Type 0x01) can be considered as a template that has a set of rules defined locally by a T-PE or S-PE. Each T-PE or S-PE can define its own set of rules and its upto the administrative domain to maintain congruities among PW community rules through which PW set-up process would follow. A LDP peer may use this community to control information it accepts, prefers or distributes to other peers.

A LDP peer receiving a PW set-up request (label mapping message) that does not carry the PW Template Community MAY append a PW Template Community TLV when propagating the label mapping message to next S-PE/T-PE.

A LDP peer receiving a PW set-up request with PW Template Community MAY modify the PW community according to local policy while propagating the request to the next-hop. Following sub-types of PW Template Community are defined in this document.

3.1.1. PW Generic Template Community

PW Generic Template Community is defined as sub-type 0x1 of the PW Template Community. The length field is 4 octets and contains a 32 bit generic identifier.




3.2. PW Color Community

A PW Color community (PW Community type 0x2) can be considered as a "coloring" of the PW that may be used by T-PE and S-PE in performing various hash functions required during PW set-up. One such application is in selection of PW signaling next-hop from multiple ECMP next-hops provided by the matching PW Route.

3.2.1. PW Generic Color Community

PW Generic Color Community is defined as sub-type 0x1 of the PW Color Community. The length field is 4 octets and contains a 32 bit generic identifier.



4. IANA Considerations

This document proposes an OPTIONAL LDP PW Communities TLV, with a proposed type of 0x407, to be allocated from the LDP TLV type registry.

5. Security Considerations

This document does not impose additional security considerations to what is defined in [[RFC5036](#)], [[RFC4447](#)] and [[I-D.ietf-pwe3-dynamic-ms-pw](#)]

6. Acknowledgements

The authors would like to acknowledge the valuable comments and suggestions from Mathew Bocci, Mustapha Aissaoui and Wim Henderickx.

7. References

7.1. Normative References

- [I-D.ietf-pwe3-dynamic-ms-pw]
Martini, L., Bocci, M., and F. Balus, "Dynamic Placement of Multi Segment Pseudowires", [draft-ietf-pwe3-dynamic-ms-pw-14](#) (work in progress), July 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", [RFC 4447](#), April 2006.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.

7.2. References

- [RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", [RFC 3107](#), May 2001.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.

Authors' Addresses

Paul Kwok
Alcatel-Lucent
701 E Middlefield Road
Mountain View, CA 94043
USA

Email: paul.kwok@alcatel-lucent.com

Pranjal Kumar Dutta
Alcatel-Lucent
701 E Middlefield Road
Mountain View, CA 94043
USA

Email: pranjal.dutta@alcatel-lucent.com

Frederic Jounay
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cidex,
France

Email: frederic.jounay@orange-ftgroup.com

