

Internet Engineering Task Force (IETF)
Request for Comments: 6667
Category: Standards Track
ISSN: 2070-1721

K. Raza
S. Boutros
C. Pignataro
Cisco Systems
July 2012

LDP 'Typed Wildcard' Forwarding Equivalence Class (FEC) for
PWid and Generalized PWid FEC Elements

Abstract

The "Typed Wildcard Forwarding Equivalence Class (FEC) Element" defines an extension to the Label Distribution Protocol (LDP) that can be used when requesting, withdrawing, or releasing all label bindings for a given FEC Element type is desired. However, a Typed Wildcard FEC Element must be individually defined for each FEC Element type. This specification defines the Typed Wildcard FEC Elements for the Pseudowire Identifier (PWid) (0x80) and Generalized PWid (0x81) FEC Element types.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6667>.

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.

This document may not be modified, and derivative works of it may not be created, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

- 1. Introduction2
- 2. Typed Wildcard for PW FEC Elements3
- 3. Applicability Statement4
- 4. Operation4
 - 4.1. PW Consistency Check5
 - 4.2. PW Graceful Shutdown5
 - 4.3. Wildcard PW Status5
 - 4.4. Typed Wildcard MAC Withdrawal in VPLS6
- 5. Security Considerations6
- 6. Acknowledgments7
- 7. References7
 - 7.1. Normative References7
 - 7.2. Informative References7

1. Introduction

An extension to the Label Distribution Protocol (LDP) [RFC5036] defines the general notion of a "Typed Wildcard Forwarding Equivalence Class (FEC) Element" [RFC5918]. This can be used when requesting, releasing, or withdrawing all label bindings for a given type of FEC Element is desired. However, a Typed Wildcard FEC Element must be individually defined for each type of FEC Element.

[RFC4447] defines the "PWid FEC Element" and "Generalized PWid FEC Element", but does not specify the Typed Wildcard format for these elements. This document specifies the format of the Typed Wildcard FEC Element for the "PWid FEC Element" and "Generalized PWid FEC Element". The procedures for Typed Wildcard processing for PWid and Generalized PWid FEC Elements are the same as described in [RFC5918] for any Typed Wildcard FEC Element type.

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 [RFC2119].

2. Typed Wildcard for PW FEC Elements

The format of the Typed Wildcard FEC Element for Pwid and Generalized Pwid is specified as:




Figure 1: Format of Typed Wildcard FEC Element for PW FEC Element Types

Where:

Typed Wcard (one octet): Typed Wildcard FEC Element type (0x05) as specified in [RFC5918].

[FEC Element] Type (one octet): PW FEC Element type:

Pwid: (type 0x80 [RFC4447])
 Generalized Pwid: (type 0x81 [RFC4447])

Len [FEC Type Info] (one octet): Two. (There is additional FEC info to scope the Typed Wildcard.)

R bit (Reserved bit): MUST be set to ZERO on transmit and ignored on receipt.

PW type (15-bits): PW type as specified in [RFC4447]. This field is used to scope the wildcard FEC operation to limit all PWs of a given type. This MUST be set to "Wildcard" type (0x7FFF), as defined in [IANA-PWE3], when referring PWs of all types (see Section 4 for its usage).

[RFC4447] defines the "PW Grouping ID TLV" that can be used for wildcard withdrawal or status messages related to Generalized Pwid FECs. When the Typed Wildcard FEC for Generalized Pwid FEC element is in use, the "PW Grouping ID TLV" MUST NOT be present in the same message. If present, the receiving Label Switching Router (LSR) MUST ignore this TLV silently and process the rest of the message.

3. Applicability Statement

The Typed Wildcard FEC Elements defined in this document for the Pwid and Generalized Pwid FEC Elements provide a finer degree of granularity when compared to the wildcard FEC mechanics defined in [RFC5036].

The Pwid FEC Element as defined in [RFC4447] contains a Group ID field. This field is defined as an arbitrary 32-bit value that represents a group of PWs and is used to create groups in the PW space, including potentially a single group of all PWs for a given FEC Element type. This grouping enables an LSR to send "wildcard" label withdrawals and/or status notification messages corresponding to a PW group upon physical port failures. Similarly, [RFC4447] defines the "PW Grouping ID TLV" used in the same fashion for the Generalized Pwid FEC Element.

The Pwid Typed Wildcard FEC Elements defined in this document help us achieve similar functionality as the "Group ID" field or "PW Grouping ID TLV" for label withdrawal and status notification messages. Additionally, the Typed Wildcard procedures [RFC5918] provide a more generalized and comprehensive solution by allowing:

1. Typed Wildcard Label Request messages
2. Label TLVs in label messages to further constrain the wildcard to all FECs of the specified FEC type [and its specific filter] that are also bound to the specified label.

This document allows use of the Typed Wildcard PW FEC Element in any LDP message that specifies a FEC TLV as a mandatory or optional parameter of the message. In addition to LDP label messages, this also applies to notification messages (containing PW Status) and Address Withdraw (for MAC address withdrawal [RFC4762]) messages in the context of LDP PW signaling. When a Typed Wildcard PW FEC Element is used in an Address Withdraw message for Virtual Private LAN Service (VPLS) Media Access Control (MAC) address withdrawal, the MAC List TLV MUST contain an empty list.

4. Operation

The use of Typed Wildcard FEC Elements for PW can be useful under several scenarios. This section describes some use cases to illustrate their application. The following use cases consider two LSR nodes, A and B, with an LDP session between them to exchange Layer 2 Virtual Private Network (L2VPN) PW bindings.

4.1. PW Consistency Check

A user may request a control-plane consistency check at LSR A for the Generalized Pwid FEC bindings that it learned from LSR B over the LDP session. To perform this consistency check, LSR A marks all its learned Generalized Pwid FEC bindings from LSR B as stale, and then sends a Label Request message towards LSR B for the Typed Wildcard FEC Element for Generalized Pwid FEC Element type with the PW type set to "Wildcard" (0x7FFF). Upon receipt of such a request, LSR B replays its database related to the Generalized Pwid FEC Element using one or more Label Mapping messages. As a PW binding is received at LSR A, the associated binding state is marked as refreshed (not stale). When replay completes for the Generalized Pwid FEC type, LSR B marks the end of its replay by sending an End-of-LIB notification [RFC5919] corresponding to the Generalized Pwid FEC Element type. Upon receipt of this notification at LSR A, any remaining stale PW binding of the Generalized Pwid FEC type learned from the peer LSR B is cleaned up and removed from the database. This completes the consistency check with LSR B at LSR A for Generalized Pwid FEC type.

4.2. PW Graceful Shutdown

It may be desirable to perform shutdown/removal of existing PW bindings advertised towards a peer in a graceful manner -- i.e., all advertised PW bindings are to be removed from a peer without session flap. For example, to request a graceful delete of the Pwid FEC and Generalized Pwid FEC bindings at LSR A learned from LSR B, LSR A would send a Label Withdraw message towards LSR B with Typed Wildcard FEC Elements pertaining to the Pwid FEC Element (with PW type set to 0x7FFF) and Generalized Pwid FEC Element (with PW type set to 0x7FFF). Upon receipt of such a message, LSR B would delete all Pwid and Generalized Pwid bindings learned from LSR A. Afterwards, LSR B would send Label Release messages corresponding to received Label Withdraw messages with the Typed FEC Element.

4.3. Wildcard PW Status

The Typed Wildcard FEC Elements for PW FECs can be very useful to convey PW status amongst LSRs. The Provider Edge (PE) devices can send the "PW Status TLV" in an LDP Notification message to indicate PW status (i.e., a Pseudowire Status Code denoting, for example, a particular fault) to their remote peers [RFC4447]. In case of a global failure affecting all PWs, an LSR typically sends one PW Status LDP Notification message per PW. This per-PW-Status message has scalability implications in a large-scale network with a large number of PWs.

Using Typed Wildcard FEC Element for a given type of PW FEC Element, the LSR will need to send only one PW Status Notification message with the Typed Wildcard PW FEC specified to notify about the common status applicable to all PWs as scoped by the PW Typed Wildcard FEC.

4.4. Typed Wildcard MAC Withdrawal in VPLS

[RFC4762] defines a pseudowire-based solution to implement Virtual Private LAN Service (VPLS). Section 6.2 of RFC 4762 describes MAC Withdrawal procedures and extensions in a VPLS environment. These procedures use the LDP Address Withdraw message containing the FEC TLV (with the PW FEC element corresponding to the VPLS instance) and MAC List TLV (to specify addresses to be withdrawn). The procedures described in RFC 4762 also allow MAC address withdrawal wildcarding for a given VPLS instance.

Using RFC 4762 procedures, a PE LSR can withdraw all MAC addresses for a given VPLS instance by sending an Address Withdraw message with a VPLS instance corresponding to the PW FEC element in a FEC TLV, and a MAC List TLV with an empty list of addresses. If there is more than one VPLS instance on a given PE LSR node, separate Address Withdraw messages need to be sent by the PE LSR if it wishes to withdraw MAC addresses for all or a subset of VPLS instances upon some global failure or configuration. Per-PW (VPLS instance) MAC Withdraw message may have some scalability implications in a large-scale network.

As stated in Section 3, this document allows use of the Typed Wildcard PW FEC in Address Withdraw messages corresponding to VPLS MAC Withdrawal. The use of PW Typed Wildcard FEC enhances the scope of MAC withdrawal beyond just a single VPLS instance and allows a PE node to wildcard withdraw all MAC addresses for:

- o all VPLS instances; or
- o all VPLS instances corresponding to a given PW type.

5. Security Considerations

No new security considerations beyond those that apply to specifications [RFC5036], [RFC4447], [RFC4762], [RFC5918], and [RFC5920] apply to the use of the PW Typed Wildcard FEC Element types described in this document.

6. Acknowledgments

The authors would like to thank Eric Rosen, Reshad Rahman, Siva Sivabalan, and Zafar Ali for their review and valuable comments. We also acknowledge Daniel Cohn for suggesting use of the Typed Wildcard PW FEC for VPLS MAC withdrawal.

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

7. References

7.1. Normative References

- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, October 2007.
- [RFC5918] Asati, R., Minei, I., and B. Thomas, "Label Distribution Protocol (LDP) 'Typed Wildcard' Forward Equivalence Class (FEC)", RFC 5918, August 2010.
- [RFC5919] Asati, R., Mohapatra, P., Chen, E., and B. Thomas, "Signaling LDP Label Advertisement Completion", RFC 5919, August 2010.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC4762] Lasserre, M., Ed., and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [IANA-PWE3] Internet Assigned Numbers Authority, "Pseudo Wires Name Spaces (PWE3)", <http://www.iana.org/assignments/pwe3-parameters>, May 2011.

Authors' Addresses

Kamran Raza
Cisco Systems, Inc.
2000 Innovation Drive
Ottawa ON K2K-3E8
Canada
EMail: skraza@cisco.com

Sami Boutros
Cisco Systems, Inc.
3750 Cisco Way
San Jose, CA 95134
USA
EMail: sboutros@cisco.com

Carlos Pignataro
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709-4987
USA
EMail: cpignata@cisco.com

