Internet Engineering Task Force (IETF)
Request for Comments: 6723
Updates: 4447, 6073
Category: Standards Track
ISSN: 2070-1721

L. Jin, Ed.
ZTE
R. Key, Ed.
Huawei
S. Delord
Alcatel-Lucent
T. Nadeau
Juniper
S. Boutros
Cisco Systems, Inc.

September 2012

Update of the Pseudowire Control-Word Negotiation Mechanism

Abstract

The control-word negotiation mechanism specified in RFC 4447 has a problem when a PE (Provider Edge) changes the preference for the use of the control word from NOT PREFERRED to PREFERRED. This document updates RFC 4447 and RFC 6073 by adding the Label Request message to resolve this control-word negotiation issue for single-segment and multi-segment pseudowires.

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/rfc6723.

Table of Contents

1. Introduction

The control-word negotiation mechanism specified in [RFC4447], Section 6.2, encounters a problem when a PE changes the preference for the use of the control word from NOT PREFERRED to PREFERRED. [RFC4447] specifies that if both endpoints prefer the use of the control word, then the pseudowire control word should be used. However, in the case where a PE changes its preference from NOT PREFERRED to PREFERRED and both ends of the PW (pseudowire) PE have the use of control word set as PREFERRED, an incorrect negotiated result of the control word as "not used" occurs. This document updates the control-word negotiation mechanism in [RFC4447] by adding a Label Request message to resolve this negotiation issue for single-segment PW. Multi-segment PW in [RFC6073] inherits the control-word negotiation mechanism in [RFC4447], and this document updates [RFC6073] by adding the processing of Label Request message on the

S-PE (Switching Provider Edge). When the PE changes the preference for the use of control word from PREFERRED to NOT PREFERRED, it should follow [RFC4447], and there is no problem.

2. 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].

3. Problem Statement

[RFC4447], Section 6, describes the control-word negotiation mechanism. Each PW endpoint has a configurable parameter that specifies whether the use of the control word is PREFERRED or NOT PREFERRED. During control-word negotiation, if one PE advertises a C bit set to 0 in the Label Mapping message with its locally configured use of control word as PREFERRED, and a corresponding peer PE changes its use of control word from NOT PREFERRED to PREFERRED, this causes an incorrect negotiated control-word result of "not used".

The following case will describe the negotiation problem in detail:

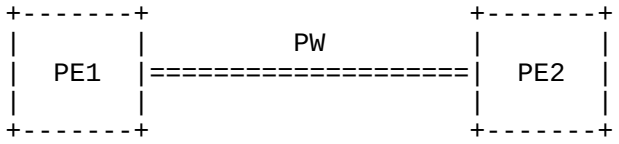

Figure 1

1. Initially, the use of control word on PE1 is configured as PREFERRED, and on PE2 as NOT PREFERRED.

2. The negotiation result for the control word of this PW is not used, and ultimately PE1 sends the Label Mapping message with C bit set to 0 according to [RFC4447], Section 6.2.

3. PE2 then changes its use of control-word configuration from NOT PREFERRED to PREFERRED, by deleting PW configuration with NOT PREFERRED use of control word, and configuring the PW again with PREFERRED use of control word.

4. PE2 will then send the Label Withdraw message to PE1, and correspondingly will receive the Label Release message from PE1.

5. According to the control-word negotiation mechanism, the previously received Label Mapping message on PE2 from PE1 carries the C bit set to 0; therefore, PE2 will still send the Label Mapping message with the C bit set to 0.

The negotiation result for the control word is still not used, even though the use of control-word configuration on both PE1 and PE2 are PREFERRED.

4. Control-Word Renegotiation by Label Request Message

The control-word negotiation mechanism in [RFC4447], Section 6, is updated to add the Label Request message described in this section.

The renegotiation process begins when the local PE has received the remote Label Mapping message with the C bit set to 0, and at the point its use of control word is changed from NOT PREFERRED to PREFERRED. The following additional procedure will be carried out:

i. The local PE MUST send a Label Release message to remote PE. If local PE has previously sent a Label Mapping message, it MUST send a Label Withdraw message to remote PE and wait until it has received a Label Release message from the remote PE. Note: the above-mentioned sending of the Label Release message and Label Withdraw message does not require a specific sequence.

ii. The local PE MUST send a Label Request message to the peer PE, and then MUST wait until it receives a Label Mapping message containing the peer's current configured preference for use of control word.

iii. After receiving the remote peer PE Label Mapping message with the C bit, the local PE MUST follow the procedures defined in [RFC4447], Section 6, when sending its Label Mapping message.

The remote PE will follow [RFC4447], and once the remote PE has successfully processed the Label Withdraw message and Label Release message, it will reset its use of control word with the locally configured preference. Then, the remote PE will send a Label Mapping message with locally configured preference for use of control word as a response to Label Request message as specified in [RFC5036].

Note: for the local PE, before processing new request to change the configuration, the above message-exchanging process should be finished. The FEC (Forwarding Equivalence Class) element in the Label Request message should be the PE's local PW FEC element. As a response to the Label Request message, the peer PE should send a

Label Mapping message with its own local PW FEC element. The Label Request message format and procedure is described in [RFC5036].

4.1. Control-Word Renegotiation for Multi-Segment PW

The multi-segment PW case for a T-PE (Terminating Provider Edge) operates similarly as the PE in single-segment PW described in the above section. An initial passive role is defined in [RFC6073] for the S-PE when processing of the Label Mapping message. [RFC6073] is updated by applying this passive role to the processing of Label Request message. When an S-PE receives a Label Request message from one of its adjacent PEs (which may be an S-PE or another T-PE), it MUST send a matching Label Request message to other adjacent PE (again, it may be an S-PE or a T-PE). This is necessary since an S-PE does not have complete information of the interface parameter field in the FEC advertisement. When the S-PE receives a Label Release message from remote PE, it MUST send a corresponding Label Release message to the other remote PE when it holds a label for the PW from the remote PE.

Note: because the local T-PE will send a Label Withdraw message before sending a Label Request message to the remote peer, the S-PE MUST process the Label Withdraw message before the Label Request message. When the S-PE receives the Label Withdraw message, it should process this message to send a Label Release message as a response and a Label Withdraw message to an upstream S-PE/T-PE. The S-PE will then process the next LDP message, e.g. the Label Request message.

When the local PE changes the use of control word from PREFERRED to NOT PREFERRED, the local PE would then renegotiate the control word so that it is not used by deleting the PW configuration with PREFERRED use of control word, and configuring the PW again with NOT PREFERRED use of control word. All of these procedures have been defined in [RFC4447], Section 5.4.1.

The diagram in Appendix A of this document updates the control-word negotiation diagram in [RFC4447] Appendix A.

4.2. Control-Word Renegotiation Use Case

The procedure of PE1 and PE2 for the use case in Figure 1 will become as follows:

1. PE2 changes locally configured preference for use of control word to PREFERRED.

2. PE2 will then send the Release messages to PE1. PE2 will also send the Label Withdraw message, and wait until it has received the Label Release message from PE1.

3. PE1 will send the Label Release message in response to the Label Withdraw message from PE2. After processing the Label Release from PE2, PE1 will then reset the use of control word to the locally configured preference as PREFERRED.

4. Upon receipt of the Label Release message from PE1, PE2 will send the Label Request message to PE1, and proceed to wait until a Label Mapping message is received.

5. PE1 will send a Label Mapping message with the C bit set to 1 again to PE2 in response to the Label Request message.

6. PE2 receives the Label Mapping message from PE1 and gets the remote label binding information. PE2 will wait for the PE1 Label Mapping message before sending its Label Mapping message with the C bit set.

7. PE2 will send the Label Mapping to PE1 with C bit set to 1, and follow procedures defined in [RFC4447], Section 6.

While it is assumed that PE1 is configured to prefer the use of the control word, in step 5, if PE1 doesn't prefer or support the control word, PE1 would then send the Label Mapping message with the C bit set to 0. As a result, PE2 in step 7 would send a Label Mapping message with the C bit set 0 as per [RFC4447], Section 6.

By sending a Label Request message, PE2 will get the locally configured preference for use of control word of peer PE1 in the received Label Mapping message. By using the new C bit from the Label Mapping message received from peer PE1 and the locally configured preference for use of control word, PE2 should determine the use of PW control word according to [RFC4447], Section 6.

5. Backward Compatibility

Since control-word negotiation mechanism is updated by adding the Label Request message, and still follows the basic procedure described in [RFC4447], Section 6, this document is fully compatible with existing implementations. For single-segment pseudowire, the remote PE (PE1 in Figure 1) which already implements [RFC4447], and the Label Request message as defined in [RFC5036] could be compatible with the PE (PE2 in Figure 1) following the mechanism of this

document. For the multi-segment pseudowire, the T-PE is the same as PE in single-segment pseudowire; the S-PE should be upgraded with the mechanism defined in this document.

6. Security Considerations

The security considerations specified in [RFC4447] and [RFC6073] also apply to this document, and this document does not introduce any additional security constraints.

7. Acknowledgements

The authors would like to thank Stewart Bryant, Andrew Malis, Nick Del Regno, Luca Martini, Venkatesan Mahalingam, Alexander Vainshtein, Adrian Farrel, and Spike Curtis for their discussion and comments.

8. Contributors

Vishwas Manral
Hewlett-Packard Co.
19111 Pruneridge Ave., Bldg. 44
Cupertino, CA 95014-0691
US
EMail: vishwas.manral@hp.com

Reshad Rahman
Cisco Systems, Inc.
2000 Innovation Drive
Ottawa, Ontario K2K 3E8
CANADA
EMail: rrahman@cisco.com

9. Normative References

[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.

[RFC6073] Martini, L., Metz, C., Nadeau, T., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", RFC 6073, January 2011.

Appendix A. Updated Diagram of C-Bit Handling Procedures

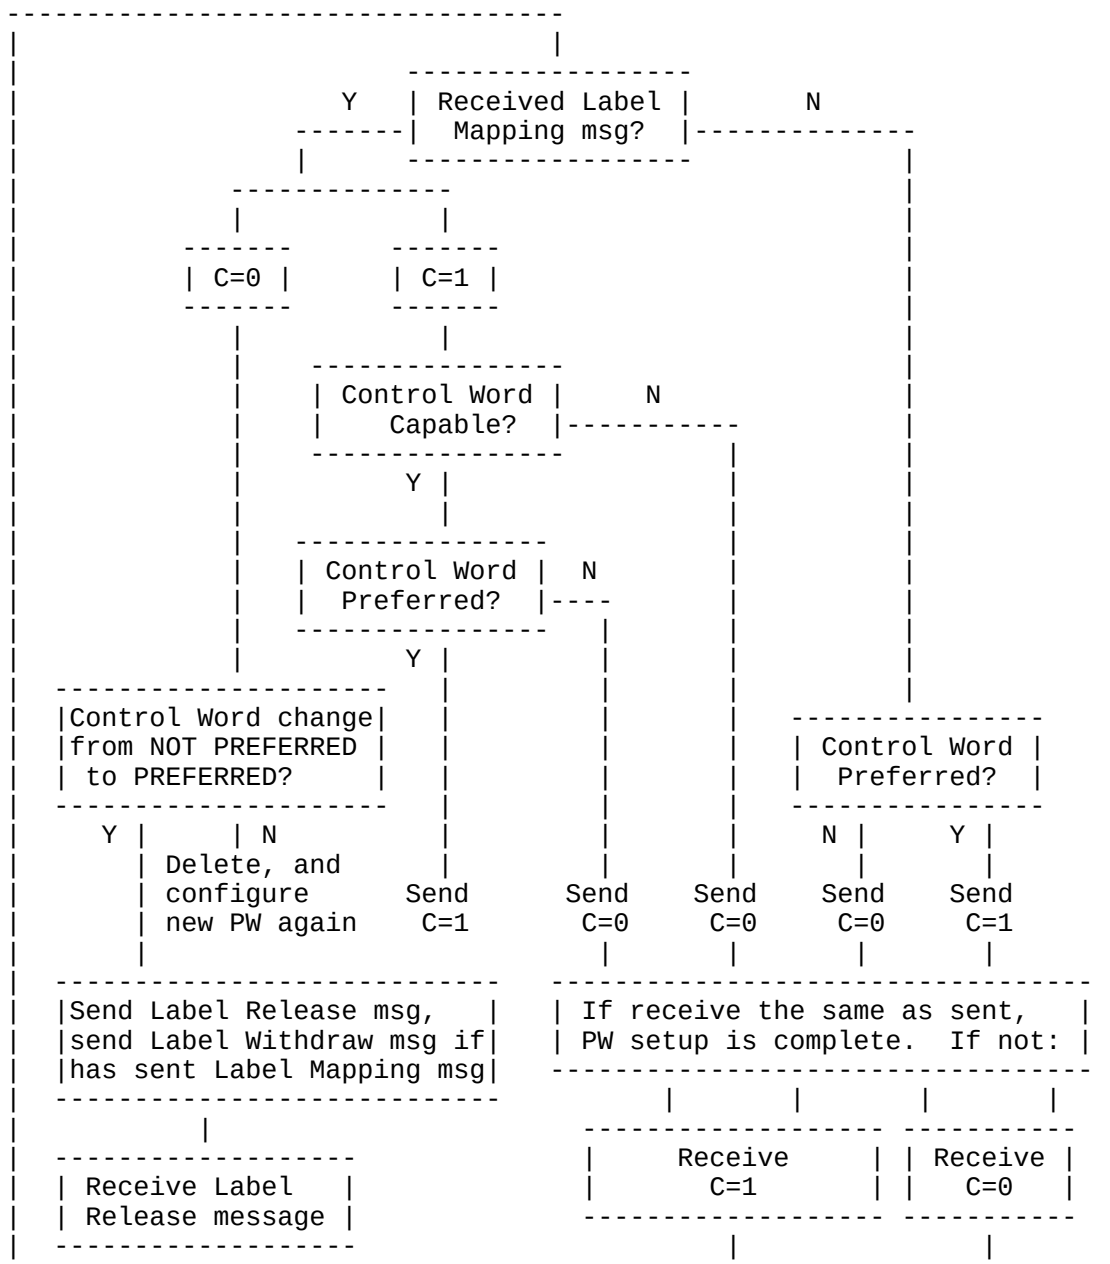

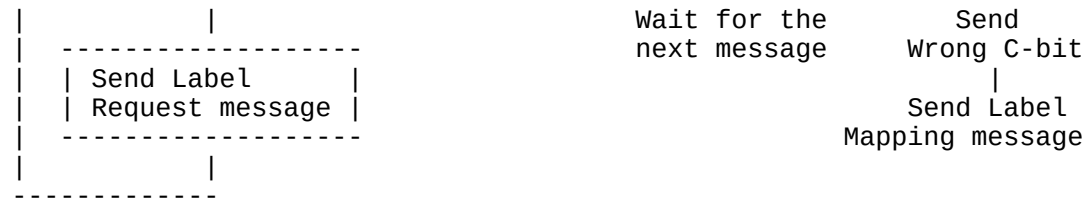

Authors' Addresses

Lizhong Jin (editor)
ZTE Corporation
889, Bibo Road
Shanghai, 201203, China

EMail: lizhong.jin@zte.com.cn

Raymond Key (editor)
Huawei

EMail: raymond.key@ieee.org

Simon Delord
Alcatel-Lucent

EMail: simon.delord@gmail.com

Thomas Nadeau
Juniper

EMail: tnadeau@juniper.net

Sami Boutros
Cisco Systems, Inc.
3750 Cisco Way
San Jose, California 95134, USA

EMail: sboutros@cisco.com