Internet Engineering Task Force (IETF)
Request for Comments: 7769
Category: Standards Track
ISSN: 2070-1721

S. Sivabalan
S. Boutros
Cisco Systems, Inc.
H. Shah
Ciena Corp.
S. Aldrin
Google Inc.
M. Venkatesan
Comcast

February 2016

Media Access Control (MAC) Address Withdrawal over Static Pseudowire

Abstract

This document specifies a mechanism to signal Media Access Control (MAC) address withdrawal notification using a pseudowire (PW) Associated Channel (ACH). Such notification is useful when statically provisioned PWs are deployed in a Virtual Private LAN Service (VPLS) or Hierarchical Virtual Private LAN Service (H-VPLS) environment.

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/rfc7769.

Table of Contents

1. Introduction

An LDP-based MAC address withdrawal mechanism is specified in [RFC4762] to remove dynamically learned MAC addresses when the source of those addresses can no longer forward traffic. This is accomplished by sending an LDP Address Withdraw Message with a MAC List TLV containing the MAC addresses to be removed from all other Provider Edge nodes over the LDP sessions. [RFC7361] describes an optimized MAC withdrawal mechanism that can be used to remove only the set of MAC addresses that need to be relearned in H-VPLS networks. [RFC7361] also describes optimized MAC withdrawal operations in PBB-VPLS networks.

A PW can be signaled via the LDP or can be statically provisioned. In the case of a static PW, an LDP-based MAC withdrawal mechanism cannot be used. This is analogous to the problem and solution described in [RFC6478] where a PW OAM (Operations, Administration, and Maintenance) message has been introduced to carry the PW status TLV using the in-band PW Associated Channel. In this document, we use a PW OAM message to withdraw MAC address(es) learned via a static PW.

Thus, MAC withdraw signaling for static PW reuses the following concepts:

- in-band signaling mechanisms used by static PW status signaling and
- MAC withdrawal mechanisms described by [RFC4762] and [RFC7361].

MAC withdraw signaling is a best effort scheme. It is an attempt to optimize network convergence by reducing blackholes caused by PW failover for protected PWs. The protocol defined in this document addresses possible loss of the MAC withdraw signal due to network congestion, but does not guarantee delivery, as is the case for the LDP-based MAC withdraw signaling. In the event that MAC withdraw signaling does not reach the intended target, the fallback to MAC re-learning due to bi-directional traffic or as a last resort aging out of MAC addresses in the absence of frames from the sources, will resume the traffic via new PW path. Such fallbacks would cause temporary blackouts but does not render a network permanently unusable.

2. Terminology

The following terminology is used in this document:

ACK: Acknowledgement for MAC withdraw message

LDP: Label Distribution Protocol

MAC: Media Access Control

MPLS: Multiprotocol Label Switching

PW: Pseudowire

PW OAM: PW Operations, Administration, and Maintenance

TLV: Type, Length, and Value

VPLS: Virtual Private LAN Services

In addition, 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. MAC Withdraw OAM Message

LDP provides reliable packet transport for control plackets for dynamic PWs. This can be contrasted with static PWs that rely on retransmission and acknowledgments (ACKs) for reliable OAM packet delivery as described in [RFC6478]. The proposed solution for MAC withdrawal over a static PW also relies on retransmissions and ACKs. However, an ACK is mandatory. A given MAC withdrawal notification is sent as a PW OAM message, and the sender retransmits the message a configured number of times in the absence of an ACK response for the sequence-numbered message. The receiver removes the MAC address(es) for a given sequence-number MAC withdraw signaling message and sends the ACK response. The receipt of the same or lower sequence-number message is responded to with an ACK but does not cause removal of MAC addresses. A new TLV to carry the sequence number has been defined.

The format of the MAC address withdraw OAM message is shown in Figure 1. The MAC withdraw PW OAM message follows the same guidelines used in [RFC6478], whereby the first 4 bytes of the OAM message header are followed by a message-specific field and a set of TLVs relevant for the message. Since the MAC withdrawal PW OAM message is not refreshed forever, a MAC address withdraw OAM message MUST contain a "Sequence Number TLV"; otherwise, the entire message is dropped. It

MAY contain the MAC Flush Parameter TLV defined in [RFC7361] when static PWs are deployed in H-VPLS and PBB-VPLS scenarios. The first 2 bits of the sequence-number TLV are reserved and MUST be set to 0 on transmit and ignored on receipt.

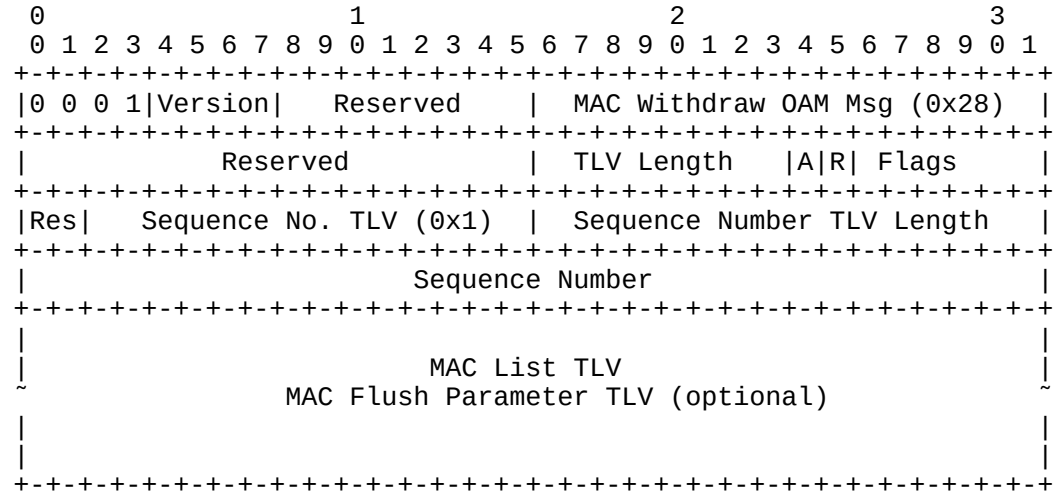

Figure 1: MAC Address Withdraw PW OAM Packet Format

In this section, the MAC List TLV and MAC Flush Parameter TLV are collectively referred to as "MAC TLV(s)". The definition and processing rules of the MAC List TLV are described by [RFC4762], and the corresponding rules of the MAC Flush Parameter TLV are governed by [RFC7361].

"TLV Length" is the total length of all TLVs in the message, and "Sequence Number TLV Length" is the length of the Sequence Number field.

A single bit (called "A-bit") is set by a receiver to acknowledge receipt and processing of a MAC Address Withdraw OAM Message. In the acknowledge message, with the A-bit set, the MAC TLVs are excluded.

A single bit (called "R-bit") is set to indicate if the sender is requesting reset of the sequence numbers. The sender sets this bit when the pseudowire is restarted and has no local record of previous send and expected receive sequence numbers.

The Sequence Number TLV MUST be the first TLV in the message.

The lack of a reliable transport protocol for the in-band OAM necessitates a presence of sequencing and acknowledgement scheme so that the receiver can recognize newer message from retransmitted older messages. [RFC4385] describes the details of sequence-number handling, which includes overflow detection for a Sequence Number field size of 16 bits. This document leverages the same scheme with the two exemptions:

- the Sequence Number field is of size 32 bits.

- overflow detection is simplified such that a sequence number that exceeds 2,147,483,647 (0x7FFFFFFF) is considered an overflow and reset to 1.

4. Operation

This section describes how the initial MAC Withdraw OAM Messages are sent and retransmitted, as well as how the messages are processed and retransmitted messages are identified.

4.1. Operation of Sender

Each PW is associated with a counter to keep track of the sequence number of the transmitted MAC withdrawal messages. Whenever a node sends a new set of MAC TLVs, it increments the transmitted sequence-number counter and includes the new sequence number in the message. The transmit sequence number is initialized to 1 at the onset, after the wrap and after the sequence number reset request receipt. Hence the transmit sequence number is set to 2 in the first MAC withdraw message sent after the sequence number is initialized to 1.

The sender expects an ACK from the receiver within a time interval we call "Retransmit Time", which can be either a default (1 second) or a configured value. If the ACK does not arrive within the Retransmit Time, the sender retransmits the message with the same sequence number as the original message. The retransmission MUST cease when an ACK is received. In order to avoid continuous retransmissions in the absence of acknowledgements, a method of suppressing retransmissions MUST be implemented. A simple and well-used approach is to cease retransmission after a small number of transmissions. In the absence of an ACK response, a one second retransmission with two retries is RECOMMENDED. However, both the interval and the number of retries are a local matter that present no interworking issues; thus, the operator MAY configure different values. Alternatively, an increasing backoff delay with a larger number of retries MAY be implemented to improve scaling issues. Whilst there are no interworking issues with any of these methods, the implementer must be mindful to not introduce network congestion and must take into

account the decaying value of the delayed MAC withdraw signaling against possible relearning due to bidirectional traffic or MAC timeout.

During the period of retransmission, if a need to send a new MAC withdraw message with updated sequence number arises, then retransmission of the older unacknowledged withdraw message MUST be suspended and retransmit time for the new sequence number MUST be initiated. In essence, a sender engages in retransmission logic only for the most recently sent withdraw message for a given PW.

In the event that a pseudowire is deleted and re-added or the router is restarted with configuration, the local node may lose information about the previously sent sequence number. This becomes problematic for the remote peer as it will continue to ignore the received MAC withdraw messages with lower sequence numbers. In such cases, it is desirable to reset the sequence numbers at both ends of the pseudowire. The reset R-bit is set in the first MAC withdraw to notify the remote peer to reset the send and receive sequence numbers. The R-bit must be cleared in subsequent MAC withdraw messages after the acknowledgement is received.

4.2. Operation of Receiver

Each PW is associated with a register to keep track of the expected sequence number of the MAC withdrawal message and is initialized to 1. Whenever a MAC withdrawal message is received, and if the sequence number on the message is greater than the value in the register, the MAC addresses contained in the MAC TLVs are removed, and the register is updated with the received sequence number. The receiver sends an ACK whose sequence number is the same as that in the received message.

If the sequence number in the received message is smaller than or equal to the value in the register, the MAC TLVs are not processed. However, an ACK with the received sequence number MUST be sent as a response. The receiver processes the ACK message as an acknowledgement for all the MAC withdraw messages sent up to the sequence number present in the ACK message and terminates retransmission.

The handling of the sequence number is described in Section 3.

A MAC withdraw message with the R-bit set MUST be processed by resetting the send and receive sequence number first. The rest of MAC withdraw message processing is performed as described above. The acknowledgement is sent with the R-bit cleared.

5. Security Consideration

The security measures described in [RFC4447], [RFC5085], and [RFC6073] are adequate for the proposed mechanism.

6. IANA Considerations

6.1. MPLS G-Ach Type

IANA has assigned a new channel type (0x0028) from the "MPLS Generalized Associated Channel (G-ACh) Types (including Pseudowire Associated Channel Types)" registry. The description of the new channel type is "MAC Withdraw OAM Message".

6.2. Sequence Number TLV

IANA has assigned a new TLV Type (0x0001) from the existing LDP "TLV Type Name Space" registry. The description for the new TLV Type is "Sequence Number TLV".

7. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <http://www.rfc-editor.org/info/rfc4385>.

[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, DOI 10.17487/RFC4447, April 2006, <http://www.rfc-editor.org/info/rfc4447>.

[RFC4762] Lasserre, M., Ed., and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, DOI 10.17487/RFC4762, January 2007, <http://www.rfc-editor.org/info/rfc4762>.

[RFC5085] Nadeau, T., Ed., and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, DOI 10.17487/RFC5085, December 2007, <http://www.rfc-editor.org/info/rfc5085>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC6073] Martini, L., Metz, C., Nadeau, T., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", RFC 6073, DOI 10.17487/RFC6073, January 2011, <http://www.rfc-editor.org/info/rfc6073>.

[RFC6478] Martini, L., Swallow, G., Heron, G., and M. Bocci, "Pseudowire Status for Static Pseudowires", RFC 6478, DOI 10.17487/RFC6478, May 2012, <http://www.rfc-editor.org/info/rfc6478>.

[RFC7361] Dutta, P., Balus, F., Stokes, O., Calvignac, G., and D. Fedyk, "LDP Extensions for Optimized MAC Address Withdrawal in a Hierarchical Virtual Private LAN Service (H-VPLS)", RFC 7361, DOI 10.17487/RFC7361, September 2014, <http://www.rfc-editor.org/info/rfc7361>.

Authors' Addresses

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario K2K 3E8
Canada

Email: msiva@cisco.com

Sami Boutros
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, CA 95134
United States

Email: sboutros@cisco.com

Himanshu Shah
Ciena Corp.
3939 North First Street
San Jose, CA 95134
United States

Email: hshah@ciena.com

Sam Aldrin
Google Inc.

Email: aldrin.ietf@gmail.com

Mannan Venkatesan
Comcast
1800 Bishops Gate Blvd
Mount Laurel, NJ 08075
United States

Email: mannan_venkatesan@cable.comcast.com