

Technical Disclosure Commons

Defensive Publications Series

March 06, 2019

DEFAULT SELECTIVE MULTICAST ROUTE TO HANDLE HETEROGENEOUS NETWORK DESIGN

Mankamana Mishra

Swadesh Agrawal

Ali Sajassi

Samir Thoria

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Mishra, Mankamana; Agrawal, Swadesh; Sajassi, Ali; and Thoria, Samir, "DEFAULT SELECTIVE MULTICAST ROUTE TO HANDLE HETEROGENEOUS NETWORK DESIGN", Technical Disclosure Commons, (March 06, 2019)
https://www.tdcommons.org/dpubs_series/2009



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

DEFAULT SELECTIVE MULTICAST ROUTE TO HANDLE HETEROGENEOUS NETWORK DESIGN

AUTHORS:

Mankamana Mishra
Swadesh Agrawal
Ali Sajassi
Samir Thoria

ABSTRACT

Techniques are described for a solution to traffic loss caused by heterogeneous network design with Ethernet Virtual Private Network (EVPN) multicast deployments. The techniques include detecting the presence of a multicast router behind a Bridge Domain and creating a default route to receive traffic from a remote source based on the detection.

DETAILED DESCRIPTION

The current widespread adoption of EVPN services, which transgresses many limitations of Virtual Private Local Area Network (LAN) Service (VPLS), introduces the need for an efficient mechanism to replicate broadcast, unknown, and multicast (BUM) traffic towards the Provider Edge devices (PEs) that participate in the same EVPN instances (EVIs).

As a simple deployment mechanism, ingress replication can be used but this would forward Multicast traffic to PEs that might not have any interested receiver, which is not efficient with respect to bandwidth and Central Processing Unit (CPU) cycles.

The IGMP (Internet Group Management Protocol) Proxy mechanism, as described in the Internet-Draft "draft-ietf-bess-evpn-igmp-mld-proxy-02" of the Internet Engineering Task Force (IETF), provides an efficient mechanism to solve the problem caused by an ingress replication mechanism and introduces Selective Multicast procedures. IGMP Proxy defines a mechanism for EVPN PEs to originate Border Gateway Protocol (BGP) route advertisements (referred to as Selective Multicast Ethernet Tag (SMET) routes) to other EVPN PEs (which are part of same EVPN instance) to show local interest and multicast traffic is forwarded on demand.

IGMP Proxy works fine with new deployments but there are some cases in which EVPN services are being deployed within existing networks in which 100% traffic loss will occur. There are two cases in which 100% traffic loss will occur. A first case in which 100% traffic loss will occur is if EVPN Layer 2 (L2) stretch is connected to a Protocol Independent Multicast (PIM) Domain. A second case in which 100% traffic loss will occur is when an IGMP Querier is positioned in a way that an EVPN PE is not the Querier.

Consider a topology as shown below in Figure 1.

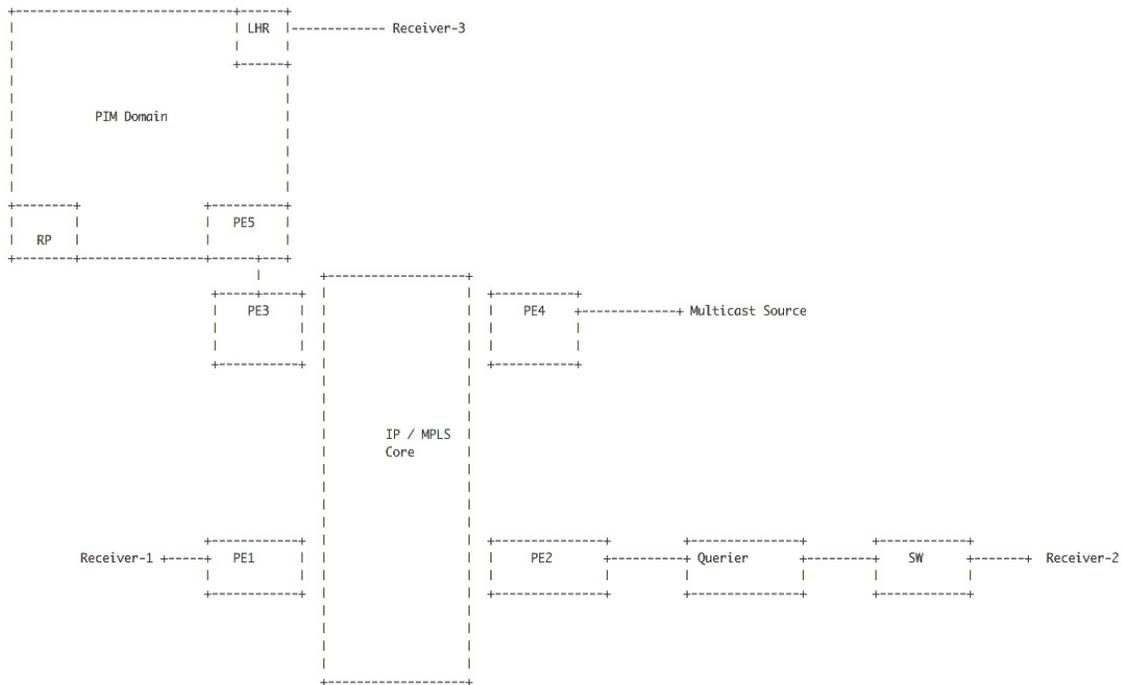


Figure 1

For the topology shown in Figure 1, PE1, PE2, PE3, and PE4 are routers which are part of an EVPN instance 1 (ESI-1). Each PE advertises an Inclusive Multicast Ethernet Tag (IMET) route with the IGMP Proxy support flag set to indicate extended community. This advertisement would imply that multicast traffic would not be forwarded before PE's originate a SMET (Selective multicast) route. Additionally for the topology, there is a multicast receiver (Receiver-1) behind PE1 which would be originating an IGMP join and also a multicast source is behind PE4. Further for the topology, there is a PIM Domain connected to a Bridge Domain behind PE3 in which the PIM domain is an Any-Source Multicast (ASM) PIM domain in which there is a Rendezvous Point (RP) in the network and a multicast receiver (Receiver-3) is behind a Last Hop Router (LHR). Finally, there is

extended Layer 2 network behind PE2. After Querier election, another device wins the Querier election other than PE2. Further down in the Layer 2 network there is multicast receiver (Receiver-2) connected to the network.

Consider the following events based on the topology shown in Figure 1. At PE4, for example, when a multicast source becomes active, PE4 receives multicast traffic. Since the rest of the other PEs in EVI-1 have already acknowledged their capability to participate in selective multicast, PE4 would not forward traffic to either of them before getting a SMET route. At PE1, when the multicast receiver (Receiver-1) sends an IGMP join, PE1 terminates the IGMP join and originates an SMET route to the other PEs where EVI-1 is active. When PE4 receives the SMET route, it starts forwarding traffic to PE1 and the multicast receiver starts getting multicast traffic.

At PE2, however, traffic loss occurs. When Receiver-2 sends the IGMP join, it would be received by the Querier and the Querier would maintain the state but the join would not be forwarded to PE2. Since PE2 does not receive any IGMP join, it does not originate any SMET route, which leads to complete traffic loss. The reason for the traffic loss is that, by nature (as per IGMP and IGMP snooping standards), the Querier expects all the traffic must pass through the Querier so that it can serve the receivers.

At PE3, traffic loss also occurs. For example, consider that Receiver-3 behind the LHR sends an IGMP join. In this case, the LHR would originate a PIM join towards the RP. However, as the RP has not yet learnt about any active source, it cannot send traffic to the receiver. Thus, PE3 does not have any information about the join request from Receiver-3, so PE3 does not originate any SMET route, and does not receive any multicast traffic.

This proposal provides a solution to the traffic losses as discussed above. In some instances, this solution can be provided for a legacy network in which there is a desire to deploy EVPN services in a manner that does not result in large modifications to the legacy network. This proposal is composed of the following components: 1) detection of a presence of a multicast router behind a Bridge Domain; and 2) creation of a default route to receive multicast traffic from a remote source.

There are several cases in which the presence of a multicast router can be detected. For example, if an EVPN enabled PE receives an IGMP Query packet from some other

device in network, this must be considered as indicating the presence of a multicast router. Another example in which a multicast router can be detected is if an EVPN enabled PE receives a PIM Hello in a Bridge Domain.

For default route creation, IGMP Proxy defines a Network Layer Reachability Information (NLRI) format to originate a SMET route, as shown below in Figure 2.

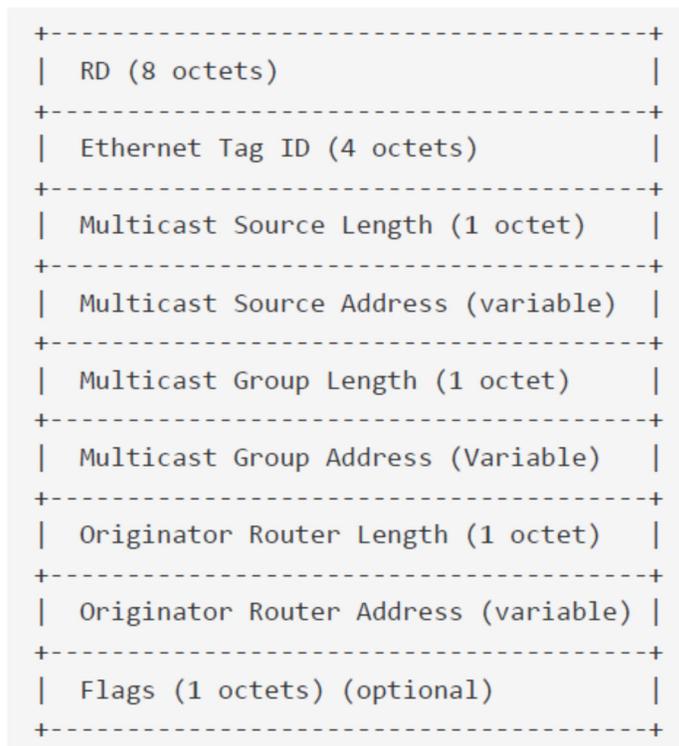


Figure 2

Upon detection of multicast router presence in a Bridge Domain, an EVPN PE must originate a route with the Multicast Source length set as zero (0) and the Multicast Group length set as zero (0). Such a change would require extension to the existing Internet-Draft "draft-ietf-bess-evpn-igmp-mld-proxy-02" to make the NLRI valid without a group address.

Consider the following events using the procedure provided by this proposal based on the topology discussed above and shown below, again, in Figure 3.

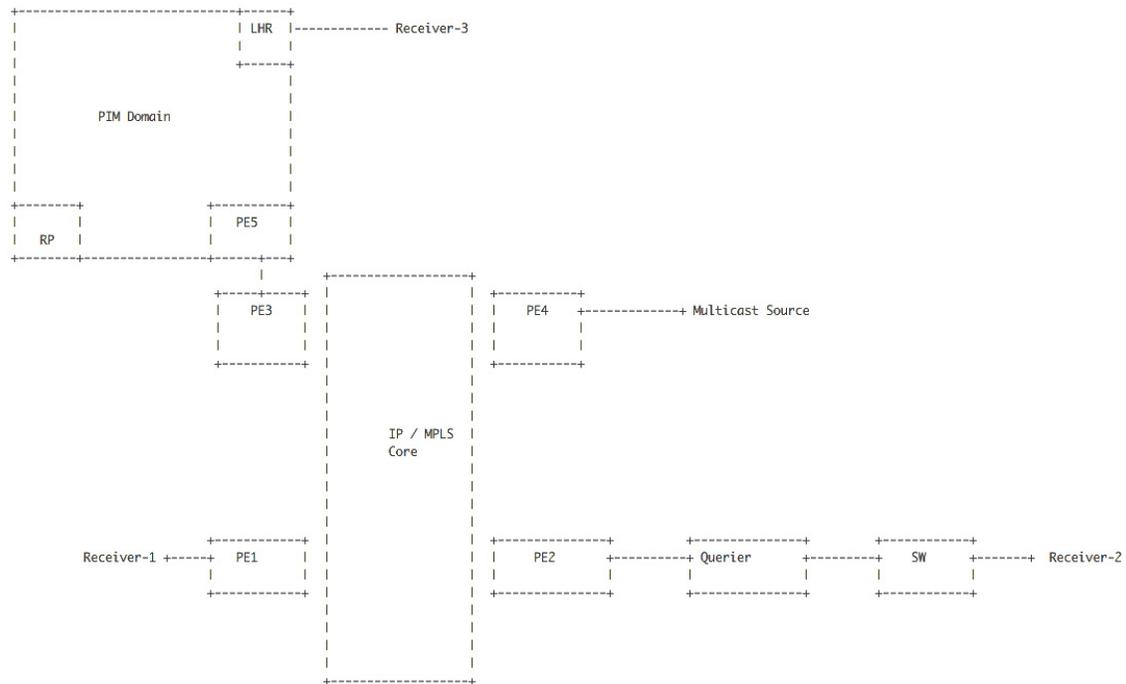


Figure 3

Using the procedure provided by this proposal at PE3, for example, since PE3 detects the PIM router behind the Bridge Domain, it originates a $(*,*)$ SMET route and expects to receive all the traffic. PE4 receives the SMET and starts forwarding traffic to PE3. PE3 would forward the traffic to mrouter port and the traffic reaches PE5. Now PE5 can perform a First Hop Router (FHR) procedure.

Regarding PE1, the events for PE1 would remain the same as previously described. At PE2 using the procedure provided by this proposal, again, PE2 detects that there is a Querier behind PE2, which makes it the mrouter port and PE2 originates a SMET route $(*,*)$. PE4 receives the SMET and starts forwarding traffic towards PE2. PE2 would forward the traffic to the Querier, which can deliver the traffic to Receiver-2. For the events at PE4, PE4 would act on the SMET routes received from remote PEs and start forwarding traffic towards the PEs.

In summary, techniques are described for detecting the presence of a multicast router behind a Bridge Domain and creating a default route to receive traffic from a remote source based on the detection, which provides a solution to traffic loss caused by heterogeneous network design with EVPN multicast deployments.