

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 16, 2013

P. Jain
G. RemaDevi
J. Kotalwar
G. Birajdar
Alcatel-Lucent, Inc.
J. Zhang
Juniper Networks Inc
May 15, 2013

BGP Extension for MVPN Site-Type Attribute
draft-jain-mvpn-bgp-sitetype-attribute-00

Abstract

This document defines a new BGP attribute in BGP based multicast VPN, that allows a MVPN PE router to inform the rest of MVPN PE routers whether it is a sender site/receiver site and there by avoid all other PEs from setting up P-tunnels to the sender site PE. This would reduce control plane states in the network and allow efficient network bandwidth utilization .

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 16, 2013.

Copyright Notice

Copyright (c) 2013 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.	Terminology	5
3.	MVPN sender site	6
4.	MVPN Site-type Attribute	7
5.	Signaling procedures and usage considerations	8
5.1.	Originating PE Procedures	8
5.2.	Receiving PE Procedures	8
6.	Management Considerations	9
7.	Security Considerations	10
8.	Acknowledgements	11
9.	IANA Considerations	12
9.1.	IANA Considerations for BGP	12
10.	References	13
10.1.	Normative References	13
10.2.	Informative References	13
	Authors' Addresses	14

Internet-Draft bgp-extensions-mvpn-site-type-attribute

May 2013

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](#) [[RFC2119](#)].

When used in lower case, these words convey their typical use in common language, and are not to be interpreted as described in [RFC2119](#) [[RFC2119](#)].

1. Introduction

[RFC 6513](#) defines sender site and receiver site in MVPN as mentioned below.

An MVPN is defined by two sets of sites: the Sender Sites set and the Receiver Sites set, with the following properties:

- o Hosts within the Sender Sites set could originate multicast traffic for receivers in the Receiver Sites set.
- o Receivers not in the Receiver Sites set should not be able to receive this traffic.
- o Hosts within the Receiver Sites set could receive multicast traffic originated by any host in the Sender Sites set
- o Hosts within the Receiver Sites set should not be able to receive multicast traffic originated by any host that is not in the Sender Sites set

With this definition the sender sites does not receive traffic but still can have terminating P-tunnels, which causes traffic in RSVP P2MP I-PMSI tunnel to be forwarded to a sender site from another sender or sender-receiver site which eventually gets dropped at the sender site. The following are the few disadvantages associated with the above approach.

- o unnecessary RSVP control plane states in the network based

- o inefficient network resource utilization.

This document addresses the above disadvantage by adding a new BGP attribute to the BGP I-PMSI A-D route which will inform other PEs that a site is sender site or receiver site, thereby preventing setting up of P-Tunnels to the sender site.

[2.](#) Terminology

Terminology used in this document:


Sender PE: PE closest to the Multicast Source (This could be either directly connected to Multicast Source or via some network).
P-Tunnel would be originating at this node. This P-Tunnel could be Inclusive or Selective..


Receiver PE: PE Node closest to the Multicast Receiver (This could be either directly connected to Multicast Source or via some network).
P-Tunnel would be terminating at this node.

Other terminologies are as defined in [[RFC6513](#)] and [[RFC6514](#)].

3. MVPN sender site

MVPN setup with sender site, receiver site and sender-receiver PEs





As seen in the above diagram, PE1 and PE2 are sender sites, PE3 is sender-receiver and PE4 is receiver site. PE1 will have terminating I-PMSI P-Tunnels from PE2 and PE3. PE2 and PE3 will send traffic towards PE1 also through these P-Tunnels. Since PE1 is a sender site, traffic received on the I-PMSI tunnels will be dropped.

The P-Tunnel which is setup to PE1 causes unnecessary control plane states in the core network.


If there is a way to inform PE2 and PE3 that PE1 is a sender site, then these PEs will not originate I-PMSI P-tunnel to PE1 and thus conserving network resources.

This document defines a new BGP attribute to be advertised to all PEs in the I-PMSI A-D route, which will inform the site-type of a PE.

4. MVPN Site-type Attribute

This document defines and uses a new BGP attribute called the "MVPN site-type attribute". This is an optional transitive BGP attribute. The format of this attribute is defined as follows:

+-----+



IANA type code (TBD)

Site Type: The field will carry the value of the site type, the value can be one of the following

- o 00 -- sender receiver site type (Default)
- o 01 -- sender site
- o 02 -- receiver site

5.1. Originating PE Procedures

A MVPN PE originating BGP I-PMSI A-D route will attach MVPN site-type attribute to the route. This attribute is used to inform the route receiving PE if the originating PE is a sender site, receiver site or a sender-receiver site

If the attribute is absent in the I-PMSI A-D route the originating PE will be considered as a sender-receiver site

If a MVPN PE with existing P-tunnels to other PEs is changed to be a sender site or receiver site, a new I-PMSI A-D route needs to be send with the new MVPN site-type attribute.

5.2. Receiving PE Procedures

A MVPN PE receiving the I-PMSI A-D route with MVPN site-type attribute performs the below action based on the value of the site-type attribute.

- o If the site-type attribute in the I-PMSI A-D route is received with value of sender site , then the route receiving PE does not include the PE which originated the I-PMSI A-D route as leaf.
- o If the site-type attribute in the I-PMSI A-D route is received with value of receiver site, then the route receiving PE includes the PE which originated the I-PMSI A-D route as leaf.
- o If the site-type attribute in the I-PMSI A-D route is received with value of sender-receiver, then the route receiving PE includes the PE which originated the I-PMSI A-D route as leaf.

If the PE receiving the BGP Intra-AD route, already has a P-Tunnel to a given PE and then it receives new I-PMSI A-D Route with site-type attribute as sender site, it must accept the I-PMSI A-D Route and tear down the existing tunnels to the sender PE.

If the PE receiving the BGP Intra-AD route has not established P-tunnel to a given PE since it is a sender site and if it receives a new I-PMSI A-D route from the PE without site-type attribute or with site-type attribute as receiver site or sender-receiver site, then it should set up a new P-tunnel to the PE.

[6.](#) Management Considerations

None

[7.](#) Security Considerations

The function described in this document does not create any new security issues for BGP protocol.

[8.](#) Acknowledgements

The authors want to thank Wim Henderickx of Alcatel-Lucent, Inc for significant contribution and feedback.

[9.](#) IANA Considerations

[9.1.](#) IANA Considerations for BGP

IANA will assign type code for MVPN site-type Attribute Flags TLV, which is carried in BGP I-PMSI Intra-AD route.

[10.](#) References

[10.1.](#) Normative References

- [RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", [RFC 6513](#), February 2012.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", [RFC 6514](#), February 2012.

[10.2.](#) Informative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1

Functional Specification", [RFC 2205](#), September 1997.

- [RFC4875] Aggarwal, R., Papadimitriou, D., and S. Yasukawa,
"Extensions to Resource Reservation Protocol - Traffic
Engineering (RSVP-TE) for Point-to-Multipoint TE Label
Switched Paths (LSPs)", [RFC 4875](#), May 2007.

Jain, et al.

Expires November 16, 2013

[Page 13]

Internet-Draft bgp-extensions-mvpn-site-type-attribute

May 2013

Authors' Addresses

Pradeep Jain
Alcatel-Lucent, Inc.
701 E Middlefield Rd
Mountain View, CA 94040
USA

Email: pradeep.jain@alcatel-lucent.com

Geetha RemaDevi
Alcatel-Lucent, Inc.
701 E Middlefield Rd
Mountain View, CA 94040
USA

Email: geetha.rema_devi@alcatel-lucent.com

Jayant Kotalwar
Alcatel-Lucent, Inc.
701 E Middlefield Rd
Mountain View, CA 94040
USA

Email: Jayant.Kotalwar@alcatel-lucent.com

Girish Birajdar
Alcatel-Lucent, Inc.
701 E Middlefield Rd
Mountain View, CA 94040
USA

Email: girish.birahdar@alcatel-lucent.com

Jeffrey Zhang
Juniper Networks Inc
10 Technology Park Drive
Westford,MA 01886
USA

Email: zzhang@juniper.net