

L3VPN  
Internet-Draft  
Intended status: Standards Track  
Expires: January 5, 2015

R. Kebler  
P. Kurapati  
Juniper Networks  
S. Asif  
AT&T LABS  
July 4, 2014

Multicast Traceroute for MVPNs  
draft-kebler-kurapati-l3vpn-mvpn-mtrace-01

## Abstract

Mtrace is a tool used to troubleshoot issues in a network deploying Multicast service. When multicast is used within a VPN service offering, the base Mtrace specification does not detect the failures. This document specifies a method of using multicast traceroute in a network offering Multicast in VPN service.

## 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 January 5, 2015.

## Copyright Notice

Copyright (c) 2014 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

Internet-Draft

MVPN Mtrace

July 2014

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

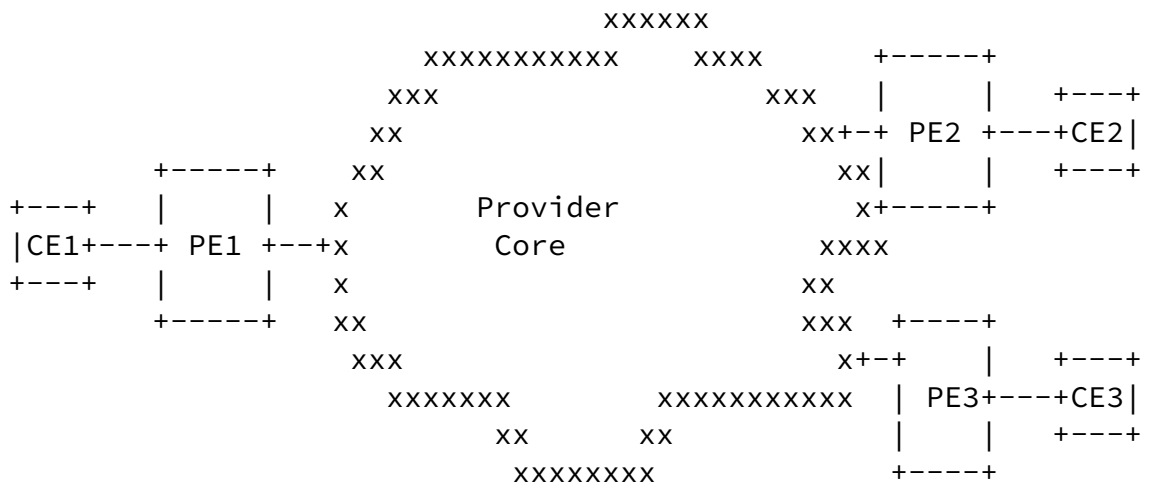
<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Overview . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Protocol Details . . . . .	<a href="#">4</a>
<a href="#">3.1.</a>	Mtrace Query . . . . .	<a href="#">4</a>
<a href="#">3.2.</a>	Mtrace Request . . . . .	<a href="#">5</a>
<a href="#">3.2.1.</a>	Ingress PE Procedures . . . . .	<a href="#">6</a>
<a href="#">3.3.</a>	Downstream Requests . . . . .	<a href="#">7</a>
<a href="#">3.4.</a>	ASBR Behavior . . . . .	<a href="#">8</a>
<a href="#">3.5.</a>	Virtual Hub and Spoke . . . . .	<a href="#">8</a>
<a href="#">3.6.</a>	Inter-Area Provider Tunnels . . . . .	<a href="#">9</a>
<a href="#">3.6.1.</a>	Egress PE . . . . .	<a href="#">9</a>
<a href="#">3.6.2.</a>	ABR Behavior . . . . .	<a href="#">9</a>
<a href="#">3.7.</a>	Mtrace MVPN Procedure . . . . .	<a href="#">9</a>
<a href="#">4.</a>	Error Detection . . . . .	<a href="#">11</a>
<a href="#">4.1.</a>	MVPN Error Codes . . . . .	<a href="#">11</a>
<a href="#">5.</a>	Mtracev2 Extensions . . . . .	<a href="#">12</a>
<a href="#">5.1.</a>	New Mtracev2 TLV Type . . . . .	<a href="#">12</a>
<a href="#">5.2.</a>	MVPN Extended Query Block . . . . .	<a href="#">12</a>
<a href="#">5.3.</a>	Leaf A-D Augmented Response Block . . . . .	<a href="#">13</a>
<a href="#">5.4.</a>	PMSI Tunnel Attributes Augmented Response Block . . . . .	<a href="#">14</a>
<a href="#">6.</a>	Mtrace2 Standard Response Block considerations . . . . .	<a href="#">14</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">14</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">15</a>
<a href="#">9.</a>	Acknowledgments . . . . .	<a href="#">15</a>
<a href="#">10.</a>	Normative References . . . . .	<a href="#">15</a>
	Authors' Addresses . . . . .	<a href="#">15</a>

## [1.](#) Introduction

The current multicast traceroute [[I-D.ietf-mboned-mtrace-v2](#)] travels up the tree hop-by-hop towards the source. This verifies the basic multicast state back to the source, but is not sufficient to verify the MVPN state. The base Mtrace specification assumes that the routers in the path are directly connected through interfaces. In the case of Multicast traffic over VPN service, the PEs who are MVPN neighbors may be separated by several router hops. The path taken by the query can be completely different from the path taken through core by the actual multicast traffic. Consider a case in the below

figure, where provider tunnel between PE2 (Source) and PE1 (Receiver) is not established correctly due to incorrect MVPN state on PE2. In the current form of Mtrace, the Query would result in a successful response since there is no error detection mechanism for MVPN state available currently. Even if one can infer from the statistics of

the Mtrace Response that PE2 has an issue, the existing error codes are not sufficient to identify the root cause. Also, there could be a problem sending traffic over the provider tunnel from PE2 to PE1, but the mtrace query will not even travel over this provider tunnel. Therefore, the mtrace successful response can be misleading. This draft ensures that the Response uses same provider-tunnel that the given C-S,C-G data would traverse and returns appropriate MVPN specific error codes which would help in identifying the root cause.



MVPN topology

## 2. Overview

As described in the Mtracev2 specification [[I-D.ietf-mboned-mtrace-v2](#)], a Querier initiates an Mtrace Query which is sent to the Last Hop Router. Last Hop Router converts this into a Request and sends it towards the First Hop Router. This draft introduces a new "Downstream Request" mechanism to allow the First Hop Router to send the mtrace request message back on the Provider

tunnel to the Last hop router. The last hop router will then change it to Response and send it to the Querier who initiated the Query. If there is any error encountered by the Last hop router or First Hop router, a Response is directly unicasted to the querier with appropriate MVPN specific error codes added. Each hop in the path of Mtrace decrements the TTL value before sending the mtrace message.

Since the Mtrace is being extended for MVPNs, the Last Hop router and First Hop router SHOULD be a Provider Edge (PE) router so that the MVPN specific error codes can be contained within the provider space. The Request will be initiated by the egress PE and will travel upstream to the ingress PE. It is assumed that the Querier knows and

can reach the egress PE. A Querier and egress PE can be the same router.

For Mtrace initiated by the CEs, the specification mentioned in Mtracev2 [[I-D.ietf-mboned-mtrace-v2](#)] SHOULD be followed. If a Mtrace message is received by the PE on CE facing interfaces containing MVPN specific extensions defined in this draft, it SHOULD be discarded.

### [3.](#) Protocol Details

The protocol details that follow are described in terms of mtracev2. However, the same procedures can be achieved with mtracev1. The protocol extensions needed for mtracev2 are described in [Section 5](#) and the protocol extensions for mtracev1 and described in [section 6](#).

#### [3.1.](#) Mtrace Query

A Querier willing to perform a Mtrace on a MVPN issues a Mtrace Query. The format of the Query TLV is as specified in the Mtracev2 specification [[I-D.ietf-mboned-mtrace-v2](#)]. The (C-S,C-G) to be queried is populated in the source address and group address fields of the Mtrace2 Query block. A deployment may use wild card SPMSIs as defined in [[RFC6625](#)]. For example, a (C-\*,C-\*) wild card SPMSI or a (C-\*,ALL-PIM-ROUTERS) can be used to send messages like BSR across PEs as mentioned in [section 5.3.4](#) MVPN specification [[RFC6513](#)]. A querier may be interested in knowing the health of such a SPMSI tunnel. In this case, the Multicast Address and Source Address fields of the Mtrace2 query can be filled with wild cards (all 1s) accordingly by the querier.

The Querier MUST add a MVPN Extended Query Block to include the RD of the C-S,C-G that it wishes to trace. When wild card SPMSIs are used, a PE could have subscribed to multiple upstream PEs for wild card SPMSIs. Hence, a query for a wild card SPMSI MUST also specify the upstream PE address that it is interested to query. The upstream PE address in the MVPN Extended Query Block MUST be filled only for wild card queries. For a regular (C-S,C-G) query, this field SHOULD be set to 0s by the querier and is ignored by the receivers.

This Query is sent to the Downstream PE (Last Hop Router) to initiate the mtrace towards the source. If a Querier does not receive a Response, it can retry sending Query messages with increasing TTL values to help diagnose where the Mtrace messages are being lost.

### [3.2.](#) Mtrace Request

The PE that receives the query will lookup the (C-S,C-G) using the RD of the query to distinguish the vrf. If the RD doesn't match any VRF, PE sends a response with error code set to BAD\_RD. The PE first checks the C-Mcast route that is matching (C-S,C-G) of the mtrace Query. It then finds the upstream multicast hop from the selected C-Mcast route and unicasts the requests to the upstream multicast hop after decrementing the TTL. The Mtrace Request MUST have PMSI Tunnel Attributes Augmented Response Block populated with the PMSI attribute that the PE uses to receive the traffic for the given (C-S,C-G) traffic.

Upstream multicast hop can be same as upstream PE router in some cases, while it can be the ASBR or the BGP nexthop of the selected C-Mcast route in Inter-AS scenarios. The procedures for finding upstream multicast hop is discussed in detail under [section 5.1](#) of MVPN specification [[RFC6513](#)].

When a wild card query is received, the PE will look for the upstream PE address in the MVPN Extended query block. The PE will then check if it has bound to the wild card SPMSI tunnel from the specified

upstream PE. If it has, it will populate the Leaf A-D Augmented Response Block and PMSI Tunnel Attributes Augmented Response Block with the respective values. If the PE has not received any wild card SPMSI AD route from the specified upstream PE in the query, it should send a response with the error code set to NO\_WILD\_CARD\_SPMSI\_AD\_RCVD. If the PE has received wild card SPMSI AD route from the upstream PE, but has not responded with a LEAF-AD route, it should send a response with the error code set to NO\_WILD\_CARD\_SPMSI\_LEAF\_AD\_SENT.

For a non-wild-card query, the upstream PE address field in the MVPN Extended query block MUST be ignored by the PEs. It MUST follow the procedure to find the upstream multicast hop as discussed earlier.

If the route does not match any MVPN-TIB state, then the PE should send a Response to the Querier with the error code set to NO\_CMCAST\_STATE. If the PE cannot locate the upstream PE then it should send a response to the Querier with the NO\_UPSTREAM\_PE error code.

From the selected UMH route, the local PE extracts the ASN of the upstream PE (as carried in the Source AS Extended Community of the route), and the source-AS field of the mtrace Query is set to that AS.

If the local and the upstream PEs are in the same AS, then the RD in the mtrace Query is set to the RD of the VPN-IP route for the source/ RP.

[Section 8](#) of MVPN specification [[RFC6513](#)] mentions two procedures (Segmented and Non-Segmented) for handling Inter-AS scenarios. If the local and the upstream PEs are in different ASes, and if segmented Inter-AS procedure is used, then the local PE finds in its VRF an Inter-AS I-PMSI A-D route whose Source AS field carries the ASN of the upstream PE. The RD of the found Inter-AS I-PMSI A-D route is used as the RD of the mtrace Query. If Inter-AS I-PMSI A-D route is not found, a response with error code UNKNOWN\_INTER\_AS is sent.

To support non-segmented inter-AS tunnels, if the local and the

upstream PEs are in different ASes, the local system finds in its VRF an Intra-AS I-PMSI A-D route from the upstream PE. The Originating Router's IP Address field of that route has the same value as the one carried in the VRF Route Import of the unicast route to the address carried in the Multicast Source field. The RD of the found Intra-AS I-PMSI A-D route is used as the RD in the mtrace Query. The Source AS field in the mtrace Query is set to value of the Originating Router's IP Address field of the found Intra-AS I-PMSI A-D route.

The PE receiving Mtrace Query will check for any errors. If any error is detected it will send the error back to the Querier. Otherwise, it will change the TLV value to be an Mtrace Request, and it will add a Mtrace2 Standard Response Block. It will also add a PMSI Tunnel Attributes Augmented Response Block with the attributes of the PMSI used to receive traffic for the S,G. If a Leaf-AD route was advertised to the upstream PE for this S,G then the PE will also include a Leaf-AD Augmented Response Block with the NLRI of the associated Leaf-AD route.

### [3.2.1.](#) Ingress PE Procedures

The PE that receives the Request, will check the PMSI attributes of the sender of request to see if they match the values used to send traffic for the S,G. If the values do not match, then the PE uses the appropriate pmsi error code as specified in 'MVPN Error Codes' section and sends a mtrace Response back to the Querier. Also, if a Leaf A-D Augmented Response Block is included, the PE will validate that it has received this Leaf A-D route from the router that sent the Request. If not, then this PE should change the error code to BAD\_LEAF\_AD and send the Response to the Querier. If the PE expects that a Leaf A-D route is needed for the downstream PE to receive traffic, but did not receive one in the mtrace Request from the sending router, then it should use a NO\_LEAF\_AD\_RCVD error code for

the mtrace Response. For a wild card SPMSI query, if the PE didn't receive LEAF AD route from the downstream PE, it should use NO\_LEAF\_AD\_RCVD error code.

When the upstream PE receives the Request, it will check for any errors. If there are errors detected, or if the TTL expired, then the PE will change the TLV code to be a Mtrace Response and unicast the response back to the Querier.

The ingress PE will also check it has local vrf connectivity for the source/RP. If it does not have any connectivity to the source/RP then it should use the base specification error code NO\_ROUTE and send an mtrace Response. Note that in a Virtual Hub and Spoke environment, it is possible for a PE to receive a mtrace Request and need to propagate it to another upstream PE. These procedures are outlined in the section "Virtual Hub and Spoke". If the PE does not expect to be receiving mtrace Responses from the mvpn core and have the route to the source located via another upstream PE, then it can use the base specification RPF\_IF error code.

If the PE that receives the Request is the ingress PE that has local vrf connectivity for the source, then it will add a Standard Response Block to the mtrace message. It will not include the additional PMSI Attributes Response Block. Then it will turn the Request into a Downstream Request by changing the value of the Type field of the TLV. It will send the mtrace message on the provider tunnel used to send the S,G data traffic.

### [3.3.](#) Downstream Requests

When a router receives Mtrace Downstream Request, it will determine if it has added any of the Response Blocks for this mtrace message. If it does not locate its address in the list of Response Blocks, then it will silently discard this mtrace message. Otherwise, it will set the 'D' bit in its PMSI Tunnel Attributes Augmented Response Block to indicate that this message has been received on the PMSI tunnel.

If this router is the egress PE that provided the initial Response Block, then it will change the mtrace type to a Reply and sends the Reply to the Querier (the egress PE and the Querier may be the same router). Otherwise, this router must send the Downstream PE on the PMSI that it would normally send traffic for the S,G. Before sending the Downstream Request, the router must decrement the TTL and check for TTL expiry. If the TTL has expired, then this router must send the Response to the Querier with the appropriate code.

### [3.4.](#) ASBR Behavior



When an ASBR receives a mtrace Request the ASBR finds an Inter-AS I-PMSI A-D route whose RD and Source AS matches the RD and Source AS carried in the mtrace Query. If no matching route is found and the ASBR is using segmented tunnels as described in MVPN specification [[RFC6513](#)], the ASBR sends an UNKNOWN\_INTER\_AS error code back to the Querier. If a matching route is found, the ASBR acts as a "first hop router" and modifies the Query type to DOWNSTREAM\_REQUEST. ASBR in this case MUST validate the PMSI attributes similar to the "first hop router" and respond if there is any errors. ASBR MUST populate PMSI Tunnel Attributes Augmented Response Block with the Inter-AS provider tunnel information before sending the DOWNSTREAM\_REQUEST. Note that the mtrace request does not proceed upstream as it is assumed that performing a traceroute and exposing IP addresses across AS boundaries would not be desirable with Segmented Inter-AS Provider Tunnels.

To support non-segmented inter-AS tunnels as described in [[RFC6513](#)], instead of matching the RD and Source AS carried in the mtrace Query against the RD and Source AS of an Inter-AS I-PMSI A-D route, the ASBR should match it against the RD and the Originating Router's IP Address of the Intra- AS I-PMSI A-D routes. The Next Hop field of the MP\_REACH\_NLRI of the found Intra-AS I-PMSI A-D route is used as the destination for the mtrace Request.

### [3.5.](#) Virtual Hub and Spoke

When a Virtual-Hub (V-HUB) as described in specification [[I-D.ietf-l3vpn-virtual-hub](#)] receives a mtrace Request the S,G may be reachable via one of its vrf interfaces. In this case, the V-HUB is an ingress PE and the procedure are defined in the Section "Ingress PE Procedures". Otherwise, the C-RP/C-S of the route is reachable via some other PE. This is the case where the received route was originated by a Virtual-spoke (V-spoke) that sees the V-HUB as the "upstream PE" for the given source, but the V-HUB sees another PE as the "upstream PE" for that source. In this case, the V-HUB should check the PMSI attributes sent in the mtrace Request against the Tunnel Attributes of the Provider Tunnel used to send traffic for the S,G from the upstream PE to the V-Spoke.

The V-HUB sends a mtrace Request to its upstream PE the same way as it would if it received a mtrace Query. V-HUB MUST add PMSI Tunnel Attributes Augmented Response Block of its own before sending the mtrace Request to the upstream PE. It may also add Leaf-AD Augmented Response Block if a Leaf-AD route was advertised upstream by the V-HUB. If the RD or Source-AS of the upstream PE is different, the V-HUB updates the MVPN Extended Query Block accordingly.

### [3.6.](#) Inter-Area Provider Tunnels

#### [3.6.1.](#) Egress PE

The egress PE does the same procedures as specified in Section "Mtrace Request" except it sends the Request upstream to the IP address determined from the Global Administrator field of the Inter-area P2MP Segmented Next-hop Extended Community as described in specification [[I-D.ietf-mpls-seamless-mcast](#)]. If the egress PE has sent a Leaf-AD route then it must send a Leaf-AD Augmented Response Block with the NLRI of the Leaf A-D route.

#### [3.6.2.](#) ABR Behavior

ABR MUST find a S-PMSI or I-PMSI route whose NLRI has the same value as the Route Key field of the received mtrace Leaf-AD extended Query Block. If such a matching route is not found then a Response should be sent to the Querier with the NO\_LEAF\_AD\_RCVD. If the ABR has sent a Leaf-AD route then it must add a Leaf-AD Augmented Response Block with the values of Leaf A-D route NLRI. The upstream node's IP address is the IP address determined from the Global Administrator field of the Inter-area P2MP Segmented Next-hop Extended Community.

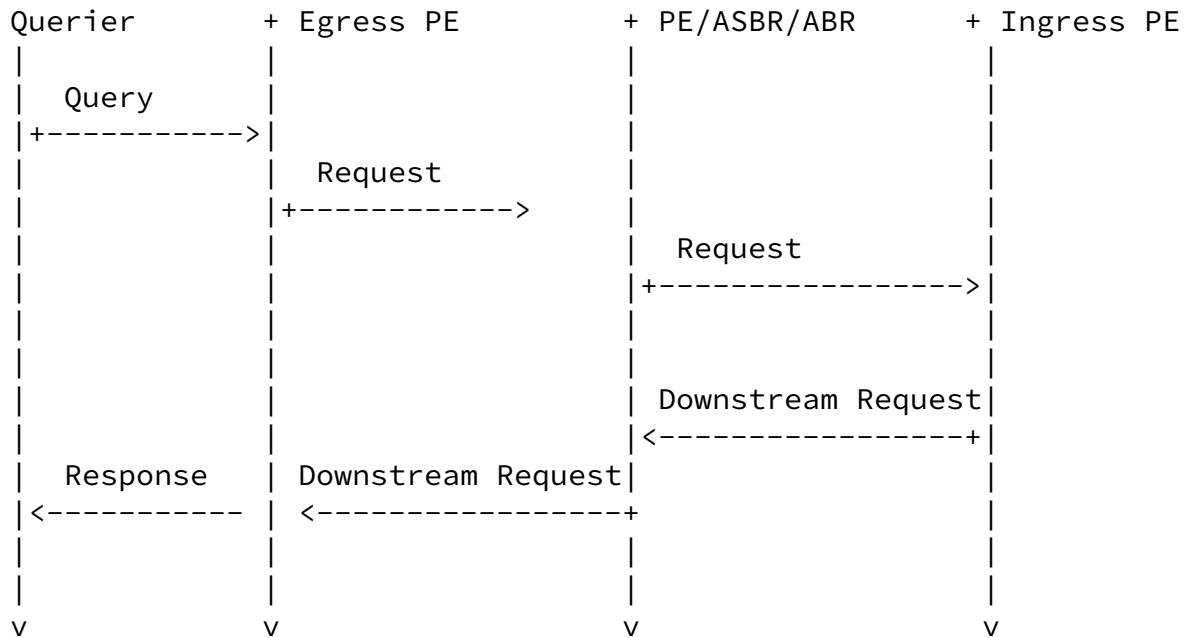
### [3.7.](#) Mtrace MVPN Procedure

In this section, we will briefly discuss the Mtrace procedure taking a working and non-working network topology.

Internet-Draft

MVPN Mtrace

July 2014



Mtrace MVPN Procedure

The above figure depicts the path of MTRACE in working condition. MTRACE request for MVPN can traverse multiple hops when a Virtual HUB is present or when segmented P2MP inter-area tunnels are used. If no error conditions are detected the downstream request will travel the same path as the regular multicast packet for the queried mroute would flow. The last hop router/egress router will convert it into a Response and send it back to querier

Let us consider a non-working case where Mtrace is expected to be used. Taking Virtual-HUB as an example, assume that there is a data-path issue between V-HUB and Egress Spoke. The below steps take place to determine the issue between V-HUB and egress Spoke

- 1 - Querier sends the Mtrace Query towards LHR (Egress PE-Spoke).
- 2 - Egress PE sends Request to V-HUB. V-HUB realises that the first hop router is a connected spoke and sends the request to Ingress Spoke PE.

3 - Ingress Spoke PE sends Downstream Request to V-HUB. The same is received by V-HUB. V-HUB sets the 'D' bit in its PMSI Tunnel Attributes Augmented Response Block.

4 - V-HUB sends Downstream request to ingress spoke. This is never received by the ingress spoke.

5 - The result of first 4 steps is that querier did not receive the response. This makes the querier fall back to TTL method.

6 - Querier reduces the TTL and the result will show that the hop from V-HUB to ingress spoke is missing thereby pointing the issue at the right place.

#### 4. Error Detection

All routers will check for normal multicast errors as defined in the Mtracev2 specification. In addition, they will check for errors specific to MVPNs and this specification.

All receiving routers will check the state of the Provider Tunnel used for forwarding traffic for the given S,G. The ability and manner to check if the Provider Tunnel is down depends on the Provider Tunnel type. If the Provider Tunnel is known to be down the PE will respond with a PTUNNEL\_DOWN error.

In some situations the router needs to send a Leaf AD route to the upstream PE. If the upstream expects a Leaf AD route, but did not receive one from the downstream PE, then the NO\_LEAF\_AD\_RCVD error will be sent.

The receiving router will check the values of the PMSI Tunnel attributes to see if they match the expected values for the PMSI. If an Inclusive-PMSI is used, then the router will verify that the values match those in the I-PMSI A-D route. If a Selective PMSI is used, then the Tunnel Attributes will be matched against the S-PMSI or Leaf A-D Route, depending on the Tunnel Type. If the values do not match, then a error code of the corresponding PMSI mismatch will be sent.

If a router receives a MVPN traceroute, but does not have the proper MVPN configuration, then it will respond with a UNEXPECTED\_MVPN error

#### [4.1.](#) MVPN Error Codes

Kebler, et al.

Expires January 5, 2015

[Page 11]

---

Internet-Draft

MVPN Mtrace

July 2014

Value	Name	Description
0x11	PTUNNEL_DOWN	The provide tunnel for this S,G i
0x12	NO_LEAF_AD_RCVD	The S-PMSI has not been joined by downstream neighbor
0x13	BAD_LEAF_AD	The Leaf A-D route does not match the expected values
0x14	BAD_RD	The RD is known to not exist on t
0x15	UNEXPECTED_MVPN	The MVPN traceroute message is un
0x16	BAD_PMSI_ATTR_FLAG	Error matching the PMSI attribute
0x17	BAD_PMSI_ATTR_TYPE	Error matching the PMSI attribute
0x18	BAD_PMSI_ATTR_LABEL	Error matching the PMSI attribute
0x19	BAD_PMSI_ATTR_ID	Error matching the PMSI attribute identifier
0x1a	UNKNOWN_INTER_AS	Could not locate the Inter-AS pro tunnel segment.

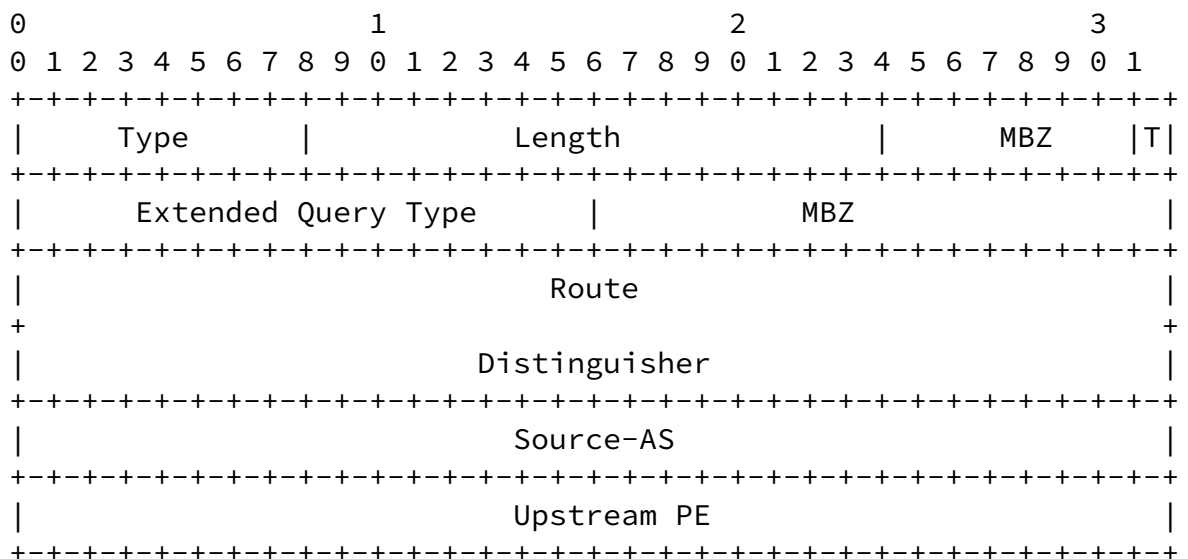
0x1b	NO_UPSTREAM_PE	No valid upstream PE or route
0x1c	NO_CMCAST_STATE	No C-Mcast route for the requester
0x1d	NO_WILD_CARD_SPMSI_AD_RCVD	No Wild Card SPMSI AD is received
0x1e	NO_WILD_CARD_SPMSI_LEAD_AD_SENT	PE did not send LEAF-AD route for the requester

## 5. Mtracev2 Extensions

### 5.1. New Mtracev2 TLV Type

A new Mtracev2 TLV type will be created for the Mtracev2 Downstream Request.

### 5.2. MVPN Extended Query Block



MVPN Extended Query Block

Type: Mtrace2 Extended Query Block Type

Length: Length of the MVPN Extended Query Block

MBZ: Sent with all 0's, ignored on receipt

T bit: This bit should be 0

Extended Query Type: New type defined


MBZ: Sent with all 0's, ignored on receipt

Route Distinguisher: The RD of the S,G that should be traced

Source-AS: The Autonomous System Number (ASN) of the Source

Upstream PE: IP Address of the Upstream PE

[5.3.](#) Leaf A-D Augmented Response Block




Leaf A-D Augmented Response Block

MBZ: Sent with all 0's, ignored on receipt

Type: New type defined

Value: The NLRI value of the associated Leaf A-D route

[5.4.](#) PMSI Tunnel Attributes Augmented Response Block



```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           Type           |D|  MBZ  | Value..  |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

### PMSI Tunnel Attributes Augmented Response Block

MBZ: Sent with all 0's, ignored on receipt

Type: New type defined

D: 'D' bit indicating that Downstream Request is received on PMSI

Value: The PMSI Tunnel Attribute as defined in [RFC 6514](#)

### 6. Mtrace2 Standard Response Block considerations

The PEs in the MVPN Mtrace add the Standard Response Block as defined in Mtrace2 [[I-D.ietf-mboned-mtrace-v2](#)]. For a PE, the incoming or outgoing interface can be a Tunnel. The First Hop Router (FHR) PE which is connected to the source SHOULD populate the incoming interface address with the respective interface connected to the CE. The outgoing interface address MAY be populated with 0 in this case. Other routers in the mtrace path MAY populate incoming and outgoing interface address fields as 0. 'Multicast Rtg Protocol' field MUST be populated with 0s by the Last Hop Router (LHR). First Hop Router (FHR) can populate this field with respective multicast routing protocol used towards its upstream CE. All the remaining fields of the Standard Response Block are populated as defined by the Mtrace2 [[I-D.ietf-mboned-mtrace-v2](#)] specification.

### 7. IANA Considerations

New TLV Type for MTRACE\_MVPN\_QUERY, MTRACE\_MVPN\_REQUEST, MTRACE\_MVPN\_DOWNSTREAM\_REQUEST, MTRACE\_MVPN\_RESPONSE

### 8. Security Considerations

There are no security considerations for this design other than what is already in the mtracev2 specification.



## 9. Acknowledgments

The authors would like to thank Yakov Rekhter and Marco Rodrigues for their valuable review and feedback.

## 10. Normative References

[I-D.ietf-l3vpn-virtual-hub]

Jeng, H., Uttaro, J., Jalil, L., Decraene, B., Rekhter, Y., and R. Aggarwal, "Virtual Hub-and-Spoke in BGP/MPLS VPNs", [draft-ietf-l3vpn-virtual-hub-08](#) (work in progress), July 2013.

[I-D.ietf-mboned-mtrace-v2]

Asaeda, H. and W. Lee, "Mtrace Version 2: Traceroute Facility for IP Multicast", [draft-ietf-mboned-mtrace-v2-10](#) (work in progress), July 2013.

[I-D.ietf-mpls-seamless-mcast]

Rekhter, Y. and R. Aggarwal, "Inter-Area P2MP Segmented LSPs", [draft-ietf-mpls-seamless-mcast-14](#) (work in progress), July 2014.

[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.

[RFC6625] Rosen, E., Rekhter, Y., Hendrickx, W., and R. Qiu, "Wildcards in Multicast VPN Auto-Discovery Routes", [RFC 6625](#), May 2012.

## Authors' Addresses

Robert Kebler  
Juniper Networks  
10 Technology Park Drive  
Westford, MA 01886  
USA

Email: [rkebler@juniper.net](mailto:rkebler@juniper.net)

Pavan Kurapati  
Juniper Networks  
1194 N. Mathilda Ave  
Sunnyvale, CA 94089  
USA

Email: kurapati@juniper.net

Saud Asif  
AT&T LABS  
200 S Laurel Ave.  
Middletown, NJ 07748  
USA

Email: sasif@att.com

