

Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 14, 2014

A. Saxena
M. Jethanandani
Ciena Corporation
October 11, 2013

Upstream mapping in Echo Request
draft-ankur-mpls-upstream-mapping-00.txt

Abstract

This document describes an enhancement to the Echo Request and Echo Response message to carry upstream mapping information for co-routed bidirectional MPLS-TP tunnels.

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 April 14, 2014.

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.

Internet-Draft

Upstream mapping

October 2013

Table of Contents

1.	Introduction	2
1.1.	Conventions Used in This Document	2
1.2.	Abbreviations	2
2.	Motivation	3
3.	Packet format	3
3.1.	Return Codes	3
3.2.	Upstream TLV	3
4.	Theory of Operations	5
4.1.	Usefulness of Upstream TLV in a Bidirectional LSP sharing the same path	5
5.	Security Considerations	6
6.	IANA Considerations	7
6.1.	New TLV	7
6.2.	New Returns Codes	7
7.	Acknowledgements	7
8.	References	7
8.1.	Normative References	7
8.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Detecting MPLS Data Plane Failures [[RFC4379](#)] defines mechanisms for collecting downstream mapping information using Downstream Mapping (DSMAP) TLV. However, it does not describe a method by which similar information can be captured for the upstream mapping. An operator would generally be interested in the path taken by a packet in both the downstream and the upstream direction. Currently the only way the operator would be able to get that information would be by running the same command from the other end point. This document describes a method by which both Downstream Mapping (DSMAP) and Upstream Mapping (UPMAP) information can be collected by the same device.

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

[1.2.](#) Abbreviations

Abbreviation	Meaning
DSMAP	Downstream Mapping

Internet-Draft

Upstream mapping

October 2013

LSP	Label Switched Path
UPMAP	Upstream Mapping

2. Motivation

Detecting MPLS Data Plane Failures [[RFC4379](#)], describes the method by which an operator can find a fault in a bidirectional LSP. The operator starts by issuing a traceroute command from a node in the network to a node that is beyond the failed node. The operator then has to issue the same command from the node that was targeted in the first command. In many cases, the operator does not have access to the other node in the network. The operator is however interested in both the upstream and downstream LSP. This draft suggests a method by which the operator can issue a single traceroute command from one of the nodes in the network and mpls echo request and response packet will carry information to validate both the DSMAP and UPMAP information. The UPMAP can only be used in case of a bidirectional LSP, where the Forward LSP and the Reverse LSP share their path. When used in a non-bidirectional LSP, the UPMAP information will be filled with zeros and SHOULD be ignored on reception. A router that does not support the UPMAP TLV will silently ignore the TLV.

3. Packet format

The packet format is similar to the packet format described in [Section 3](#) of RFC4379. [[RFC4379](#)]

This draft proposes to add two new return codes as outlined in section and a new TLV as specified in section .

3.1. Return Codes

Value	Meaning
TBD	Upstream Mapping Mismatch
TBD	Downstream and Upstream Mapping Mismatch

3.2. Upstream TLV

The upstream mapping TLV is an object that MUST be included for all reply modes in the MPLS Echo packet when the operator has requested a traceroute on a bidirectional LSP, where the Forward LSP and Reverse LSP share the same path. The presence of an upstream TLV by the requester means that the replying router SHOULD validate the upstream TLV and if correct, fill the upstream TLV with upstream FEC of the replying router. If incorrect, it should fill the return code with one of the values specified in section to indicate "Upstream Mapping Mismatch" and leave the upstream TLV as is. If the node is an LER router and the upstream TLV is included in the MPLS echo request packet, it SHOULD fill the upstream TLV with the appropriate information and MUST include it in the MPLS echo reply.

As defined in [RFC 4379](#), the length of this TLV is $K + M + 4 \times N$ octets, where M is the Multipath Length, and N is the number of Downstream Labels. Values for K are found in the description of Address Type below. The Value field of a Upstream TLV has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
MTU										Address Type										DS Flags																			
Upstream IP Address (4 or 16 octets)																																							
Upstream Interface Address (4 or 16 octets)																																							
Multipath Type										Depth Limit										Multipath Length																			

[illegible]

Upstream IP Address and Upstream Interface Address

IPv4 addresses and interface indices are encoded in 4 octets; IPv6 addresses are encoded in 16 octets. If the interface to the upstream node is numbered, then the Address Type MUST be set to IPv4 or IPv6,

the Upstream IP Address MUST be set to either the Upstream node's Router ID or the interface address of the Upstream node, and the Upstream Interface Address MUST be set to the upstream node's interface address. If the interface to the upstream node is unnumbered, the Address Type MUST be IPv4 Unnumbered or IPv6 Unnumbered, the Upstream IP Address MUST be the upstream node's Router ID, and the Upstream Interface Address MUST be set to the index assigned by the node to the interface.

If a node does not know the IP address of its neighbor, then it **MUST** set the Address Type to either IPv4 Unnumbered or IPv6 Unnumbered. For IPv4, it must set the Upstream IP Address to 127.0.0.1; for IPv6 the address is set to 0::1. In both cases, the interface index **MUST** be set to 0.

Upstream Label(s)

The set of labels in the label stack should appear as if this router were forwarding the packet through this interface. Any Implicit Null labels are explicitly included. Labels are treated as numbers, i.e., they are right justified in the field.

A Upstream Label is 24 bits, in the same format as an MPLS label minus the TTL field, i.e., the MSBit of the label is bit 0, the LSBit is bit 19, the EXP bits are bits 20-22, and bit 23 is the S bit. The replying router SHOULD fill in the EXP and S bits; the LSR receiving the echo reply MAY choose to ignore these bits.



For explanation of rest of the fields in the Upstream TLV please refer [section 3.3](#) of Detecting MPLS Data Plane Failures [[RFC4379](#)].

4. Theory of Operations

4.1. Usefulness of Upstream TLV in a Bidirectional LSP sharing the same path

The Upstream TLV MUST only be used in case of a bidirectional LSP where Forward and Reverse Paths are same, for example, MPLS-TP Co-routed tunnels or Multisegment Pseudo wire. In which case, the transit nodes will know all the information required to fill both the Downstream Mapping TLV and Upstream TLV.

Consider the following example:



In the above fig, LER1 is the ingress node with forward out going label L0 and reverse in coming label of L3. LSR1 is the transit router with forward incoming and outgoing labels as L0 and L1 respectively and reverse incoming and outgoing labels of L2 and L3 respectively. LER2 is the egress router with forward incoming label of L1 and reverse outgoing label of L2.

The ingress node SHOULD fill its Downstream TLV for label L0 and Upstream TLV for label L3. When this MPLS Echo request packet (containing the Upstream TLV and the DownStream TLV) reaches the transit node, then the node validates both Upstream TLV for label L3 and Downstream TLV for Label L0. If the Downstream TLV for label L0

specified in the packet does not match the information the transit node has, then the transit node sends a return code specifying Downstream TLV mismatch. Similarly, if the Upstream TLV specified in the packet does not match the Upstream information the transit node has, then the transit node SHOULD send a return code of Upstream TLV mismatch. If both, the Upstream TLV and Downstream TLV does not match then the transit node should send a return code of Upstream and Downstream TLV mismatch. And if both the TLVs match then the transit node populates it's Downstream Mapping for label L1 and the Upstream Mapping for label L2 and sends the reply back to the ingress node. The ingress node uses this new Downstream TLV and Upstream TLV in it's next Echo Request packet. The egress node on receiving the Echo Request packet validates Upstream TLV and Downstream TLV. If both the TLVs match then the egress node SHOULD send a return code of Replying router is egress, else it SHOULD send the return code depending on which TLV did not match.

In case a bidirectional LSP does not share the Forward and Reverse path, for example, MPLS-TP Associated LSPs, traceroute SHOULD NOT add Upstream TLV as part of the MPLS Echo Request. If the Forward and Reverse LSPs are not on the same node then the transit node of the Forward LSP won't have any information to fill the Upstream TLV.

5. Security Considerations

Security considerations, as discussed in Detecting MPLS Data Plane Failures [[RFC4379](#)], are applicable to this document.

6. IANA Considerations

6.1. New TLV

IANA would have to assign a new TLV value to the following TLV from the "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry, "TLVs and sub-TLVs" sub-registry.

Upstream Detailed Mapping TLV (see Section).

[6.2.](#) New Returns Codes

IANA needs to assign a new Return Code values from the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry, "Return Codes" sub-registry, as follows using a Standards Action value.

+-----+-----+-----+-----+	
Value	Meaning
+-----+-----+-----+-----+	
TBD	Upstream mapping mismatch
TBD	Downstream and Upstream mapping mismatch
+-----+-----+-----+-----+	

[7.](#) Acknowledgements

We would like to thank Ashesh Mishra and Vijay D'Souza for their feedback on this draft.

[8.](#) References

[8.1.](#) Normative References

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), February 2006.

[8.2.](#) Informative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC6424] Bahadur, N., Kompella, K., and G. Swallow, "Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels", [RFC 6424](#), November 2011.

Authors' Addresses

Ankur Saxena
Ciena Corporation
3939 N. First Street
San Jose, CA 95134
USA

Phone: +1 (408) 904-2109
Email: ankurpsaxena@gmail.com

Mahesh Jethanandani
Ciena Corporation
3939 N. First Street
San Jose, CA 95134
USA

Phone: +1 (408) 904-2160
Email: mjethanandani@gmail.com