

Internet Engineering Task Force
Internet-Draft
Updates: 4379 (if approved)
Intended status: Standards Track
Expires: June 18, 2014

N. Akiya
G. Swallow
C. Pignataro
Cisco Systems
L. Andersson
M. Chen
Huawei
December 15, 2013

Label Switched Path (LSP) Ping/Trace Reply Mode Simplification
draft-akiya-mpls-lsp-ping-reply-mode-simple-01

Abstract

This document adds one reply mode to indicate reverse LSP, to be used by Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Ping and Traceroute. This document also adds an optional TLV which can carry ordered list of reply modes.

This document updates [RFC4379].

Requirements Language

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 RFC 2119 [RFC2119].

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 June 18, 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.

Table of Contents

1. Introduction	2
2. Problem Statements	3
3. Solution	4
3.1. Reply via reverse LSP	4
3.2. Reply Mode Order TLV	5
4. Security Considerations	6
5. IANA Considerations	6
5.1. New Reply Mode	6
5.2. New Reply Mode Order TLV	6
6. Acknowledgements	7
7. Contributing Authors	7
8. References	7
8.1. Normative References	7
8.2. Informative References	7
Authors' Addresses	7

1. Introduction

MPLS LSP Ping, described in [RFC4379], allows initiator to encode instructions (Reply Mode) on how responder is to send response back to the initiator. [I-D.ietf-mpls-return-path-specified-lsp-ping] also allows initiator to encode a TLV (Reply Path TLV) which can instruct responder to use specific LSP to send response back to the initiator. Both approaches are powerful as they provide ability for the initiator to control the return path.

It is, however, becoming increasingly difficult for an initiator to select the "right" return path to encode in MPLS LSP echo request packets. Consequence of initiator not selecting the "right" return path encoding can result in false failure of MPLS LSP Ping and Traceroute operations, due to initiator not receiving back expected MPLS LSP echo reply. Resulting from an effort to minimize such false failures, implementations may result in having different "default"

return path encoding per LSP type and per operational type. Deviating "default" return path encoding, potentially, per vendor per LSP type per operational type can drift this technology from consistency axle. Thus it is desirable to have a return path encoding mechanism which minimizes false failure scenarios while reverse direction taking path preferred for operational needs.

2. Problem Statements

It is becoming increasingly difficult for implementations to automatically supply a workable return path encoding for all MPLS LSP Ping and Traceroute operations across all LSP types. There are several factors which are contributing to this complication.

- o Some LSPs have control-channel, and some do not. Some LSPs have reverse LSP, and some do not. Some LSPs have IP route in reverse direction, and some do not.
- o LSRs on some LSPs can have different available return path(s). Available return path(s) can depend on whether responder is a transit LSR or an egress LSR. In case of bi-directional LSP, available return path(s) on transit LSRs can also depend on whether LSP is completely co-routed, partially co-routed or non-co-routed.
- o MPLS LSP echo request packets may falsely terminate on an unintended target which can have different available return path(s) than intended target.
- o MPLS LSP Ping operation is expected to terminate on egress LSR. However, MPLS LSP Ping operation with specific TTL values and MPLS LSP Traceroute operation can terminate on both transit LSR(s) and egress LSR.

Except for the case where responder node does not have an IP route back to the initiator, it is possible to use Reply Mode of value 2 (Reply via an IPv4/IPv6 UDP packet) in all cases. However, some operators are preferring control-channel and reverse LSP as "default" return path if they are available, which are not always available.

When specific return path encoding is being supplied by users or applications, then there are no issues in choosing the return path encoding. When specific return path encoding is not being supplied by users or applications, then implementations require extended logic to compute, and sometimes "guess", the "default" return path encodings. If a responder received a MPLS LSP echo request containing return path instruction which cannot be accommodated due to unavailability, then responder implementations often drop such

packets. This results in initiator to not receive back MPLS LSP echo reply packets. Consequence may be acceptable for failure cases (ex: broken LSP) where MPLS LSP echo request terminated on unintended target. However, initiator not receiving back MPLS LSP echo reply packets, even when intended target received and verified the requests, is not desirable as result will be conveyed as false failures to users.

Some return path(s) are more preferred than others, but preferred cannot be used in all cases. Thus implementations are required to compute when preferred return path encoding can and cannot be used, and that computation is becoming more and more difficult.

This document adds one Reply Mode to describe reverse LSP, and one optional TLV to describe ordered list of reply modes. Based on operational needs, the TLV can describe multiple Reply Mode values in preferred order to allow responder to use first available Reply Mode from the list. This eliminates the need for initiator to compute, or sometimes "guess", the "default" return path encoding. And that will result in simplified implementations across vendors, and result in improved usability to fit operational needs.

3. Solution

This document adds one reply mode to indicate reverse LSP, to be used by Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Ping and Traceroute. This document also adds an optional TLV which can carry ordered list of reply modes.

3.1. Reply via reverse LSP

Some LSP types are capable of having related LSP in reverse direction, through signaling or other association mechanisms. This document uses the term "Reverse LSP" to refer to the LSP in reverse direction of such LSP types. Note that this document isolates the scope of "Reverse LSP" applicability to those reverse LSPs which are capable of and permitted to carry the IP encapsulated MPLS LSP echo reply.

This document adds one Reply Mode to be used by MPLS LSP Ping and Traceroute operations.

Value	Meaning
-----	-----
TBD1	Reply via reverse LSP

MPLS LSP echo request with TBD1 (Reply via reverse LSP) in the Reply Mode field may be used to instruct responder to use reverse LSP to

send MPLS LSP echo reply. Reverse LSP is in relation to the last FEC specified in the Target FEC Stack TLV.

When responder is using this Reply Mode, transmitting MPLS LSP echo reply packet MUST use IP destination address of 127/8 for IPv4 and 0:0:0:0:0:FFFF:7F00/104 for IPv6.

3.2. Reply Mode Order TLV

This document also introduces a new optional TLV to describe list of Reply Mode values. The new TLV will contain one or more Reply Mode value(s) in preferred order, first Reply Mode value appearing being most preferred. Following rules apply when using Reply Mode Order TLV.

1. Reply Mode Order TLV MAY be included in MPLS echo request.
2. Reply Mode Order TLV MUST NOT be included in MPLS echo reply.
3. Reply Mode field of MPLS echo request MUST be set to a valid value when supplying Reply Mode Order TLV in transmitting MPLS echo request. It is RECOMMENDED for initiator to set Reply Mode field of MPLS echo request to a value that corresponds to return path most likely to be available, in case responder does not understand the Reply Mode Order TLV.
4. If responder node understands Reply Mode Order TLV and TLV is valid, then responder MUST consider Reply Mode values described in the TLV and MUST NOT use value described in Reply Mode field of received MPLS echo request.
5. If responder node understands Reply Mode Order TLV but TLV is not valid (due to conditions listed below), then responder MUST only use value described in Reply Mode field of received MPLS echo request.
6. Reply Mode Order TLV MUST contain at least one Reply Mode value, and SHOULD contain at least two Reply Mode values.
7. Same Reply Mode value MUST NOT appear multiple times in the Reply Mode Order TLV.
8. Reply Mode value 1 (Do not reply) SHOULD NOT be used in the Reply Mode Order TLV.

The responding node is to select the first available return path in this TLV. Reply Mode value corresponding to selected return path MUST be set in Reply Mode field of MPLS LSP echo reply to communicate back to the initiator which return path was chosen.

The format of the TLV is as follows:




Figure 1 Reply Mode Order TLV

This is a variable length optional TLV. Each Reply Mode field is 1 octet.

4. Security Considerations

Beyond those specified in [RFC4379], there are no further security measures required.

5. IANA Considerations

5.1. New Reply Mode

IANA is requested to assign one reply modes from the "Reply Mode" sub-registry within the "Multiprotocol Label Switching Architecture (MPLS)" registry.

Value	Meaning	Reference
-----	-----	-----
TBD1	Reply via reverse LSP	this document

5.2. New Reply Mode Order TLV

IANA is requested to assign a new TLV type value from the "TLVs" sub-registry within the "Multiprotocol Label Switching Architecture (MPLS)" registry, for the "Reply Mode Order TLV".

The new TLV Type value should be assigned from the range (32768-49161) specified in RFC 4379 [RFC4379] section 3 that allows the TLV type to be silently dropped if not recognized.

Type	Meaning	Reference
----	-----	-----
TBD2	Reply Mode Order TLV	this document

6. Acknowledgements

Authors would like to thank Santiago Alvarez and Faisal Iqbal for discussions which motivated creation of this document. Authors would also like to thank Sam Aldrin and Curtis Villamizar for providing valuable comments to influence the contents of the draft.

7. Contributing Authors

Shaleen Saxena
Cisco Systems
Email: ssaxena@cisco.com

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

8.2. Informative References

- [I-D.ietf-mpls-return-path-specified-lsp-ping]
Chen, M., Cao, W., Ning, S., JOUNAY, F., and S. DeLord,
"Return Path Specified LSP Ping", draft-ietf-mpls-return-
path-specified-lsp-ping-15 (work in progress), October
2013.

Authors' Addresses

Nobo Akiya
Cisco Systems

Email: nobo@cisco.com

George Swallow
Cisco Systems

Email: swallow@cisco.com

Carlos Pignataro
Cisco Systems

Email: cpignata@cisco.com

Loa Andersson
Huawei

Email: loa@mail01.huawei.com

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com