Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 17, 2013

M. Chen
W. Cao
Huawei Technologies Co., Ltd
S. Ning
Tata Communications
F. Jounay
Orange CH
S. Delord
Alcatel-Lucent
September 13, 2012

Return Path Specified LSP Ping
draft-ietf-mpls-return-path-specified-lsp-ping-10.txt

Abstract

This document defines extensions to the failure-detection protocol for Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) known as "LSP Ping" that allow selection of the LSP to use for the echo reply return path. Enforcing a specific return path can be used to verify bidirectional connectivity and also increase LSP ping robustness. It may also be used by Bidirectional Forwarding Detection (BFD) for MPLS bootstrap signaling thereby making BFD for MPLS more robust.

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 March 17, 2013.

Table of Contents

1. Introduction

This document defines extensions to the failure-detection protocol for Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) known as "LSP Ping" [RFC4379] that can be used to specify the return paths for the echo reply message, increasing the robustness of LSP Ping, reducing the opportunity for error, and improving the reliability of the echo reply message. A new reply mode, which is referred to as "Reply via Specified Path", is added and a new Type-Length-Value (TLV), which is referred to as Reply Path (RP) TLV, is defined in this memo.

With the extensions described in this document, a bidirectional LSP and a pair of unidirectional LSPs (one for each direction) could both be tested with a single operational action, hence providing better control plane scalability. The defined extensions can also be utilized for creating a single Bidirectional Forwarding Detection (BFD)[RFC5880], [RFC5884]session for a bidirectional LSP or for a pair of unidirectional LSPs (one for each direction).

In this document, term bidirectional LSP includes the co-routed bidirectional LSP defined in [RFC3945] and the associated bidirectional LSP that is constructed from a pair of unidirectional LSPs (one for each direction), and which are associated with one another at the LSP's ingress/egress points [RFC5654]. The mechanisms defined in this document can apply to both IP/MPLS and MPLS Transport Profile (MPLS-TP) scenarios.

2. Problem Statements and Solution Overview

MPLS LSP Ping is defined in [RFC4379]. It can be used to detect data path failures in all MPLS LSPs, and was originally designed for unidirectional LSPs.

LSP are increasingly being deployed to provide bidirectional services. The co-routed bidirectional LSP is defined in [RFC3471] and [RFC3473], and the associated bidirectional LSP is defined in [RFC5654]. With the deployment of such services, operators have a desire to test both directions of a bidirectional LSP in a single operation.

Additionally, when testing a single direction of an LSP (either a unidirectional LSP, or a single direction of a bidirectional LSP) using LSP Ping, the validity of the result may be affected by the success of delivering the echo reply message. Failure to exchange these messages between the egress Label Switching Router (LSR) and the ingress LSR can lead to false negatives where the LSP under test

is reported as "down" even though it is functioning correctly.

2.1. Limitations of Existing Mechanisms for Bidirectional LSPs

With the existing LSP Ping mechanisms as defined in [RFC4379], operators have to enable LSP detection on each of the two ends of a bidirectional LSP independently. This not only doubles the workload for the operators, but may also bring additional difficulties when checking the backward direction of the LSP under the following conditions:

1. The LSR that the operator logged on to perform the checking operations might not have out-of-band connectivity to the LSR at the far end of the LSP. That can mean it is not possible to check the return direction of a bidirectional LSP in a single operation - the operator must log on to the LSR at the other end of the LSP to test the return direction.

2. The LSP being tested might be an inter-domain/inter-AS LSP where the operator of one domain/AS may have no right to log on to the LSR at the other end of the LSP since this LSR resides in another domain/AS. That can make it completely impossible for the operator to check the return direction of a bidirectional LSP.

Associated bidirectional LSPs have the same issues as those listed for co-routed bidirectional LSPs.

This document defines a mechanism to allow the operator to request that both directions of a bidirectional LSP be tested by a single LSP Ping message exchange.

2.2. Limitations of Existing Mechanisms for Handling Unreliable Return Paths

[RFC4379] defines 4 reply modes:

1. Do not reply
2. Reply via an IPv4/IPv6 UDP packet
3. Reply via an IPv4/IPv6 UDP packet with Router Alert
4. Reply via application level control channel.

Obviously, the issue of the reliability of the return path for an echo reply message does not apply in the first of these cases.

[RFC4379] states that the third mode may be used when the IP return path is deemed unreliable. This mode of operation requires that all intermediate nodes must support the Router Alert option and must

understand and know how to forward MPLS echo replies.

This is a rigorous requirement in deployed IP/MPLS networks especially since the return path may be through legacy IP-only routers. Furthermore, for inter-domain LSPs, the use of the Router Alert option may encounter significant issues at domain boundaries where the option is usually stripped from all packets. Thus, the use of this mode may itself introduce issues that lead to the echo reply messages not being delivered.

And in any case, the use modes 2 or 3 cannot guarantee the delivery of echo responses through an IP network that is fundamentally unreliable. The failure to deliver echo response messages can lead to false negatives making it appear that the LSP has failed.

Allowing the ingress LSR to control the path used for echo reply messages, and in particular forcing those messages to use an LSP rather than being sent through the IP network, enables an operator to apply an extra level of deterministic process to the LSP Ping test.

This document defines extensions to LSP Ping that can be used to specify the return paths of the echo reply message in an LSP echo request message.

3. Extensions

LSP Ping defined in [RFC4379] is carried out by sending an echo request message. It carries the Forwarding Equivalence Class (FEC) information of the LSP being tested which indicates which MPLS path is being verified, along the same data path as other normal data packets belonging to the FEC.

LSP Ping [RFC4379] defines four reply modes that are used to direct the egress LSR in how to send back an echo reply. This document defines a new reply mode, the Reply via Specified Path mode. This new mode is used to direct the egress LSR of the tested LSP to send the echo reply message back along the path specified in the echo request message.

In addition, two new TLVs, the Reply Path TLV and Reply Traffic Class (TC) [RFC5462] TLV, are defined in this document. The Reply Path TLV contains one nested sub-TLV that can be used to carry the specified return path information to be used by the echo reply message.

3.1. Reply Via Specified Path mode

A new reply mode is defined to be carried in the Reply Mode field of the LSP Ping echo request message.

The recommended value of the Reply via Specified Path mode is 5 (This is to be confirmed by the IANA).

Value	Meaning
5	Reply via Specified Path

The Reply via Specified Path mode is used to request that the remote LSR receiving the LSP Ping echo request message sends back the echo reply message along the specified paths carried in the Reply Path TLV.

3.2. Reply Path (RP) TLV

The Reply Path (RP) TLV is an optional TLV, if the Reply via Specified Path mode requested, the Reply Path (RP) TLV MUST be included in an echo request message. It carries the specified return paths that the echo reply message is required to follow. The format of Reply Path TLV is as follows:

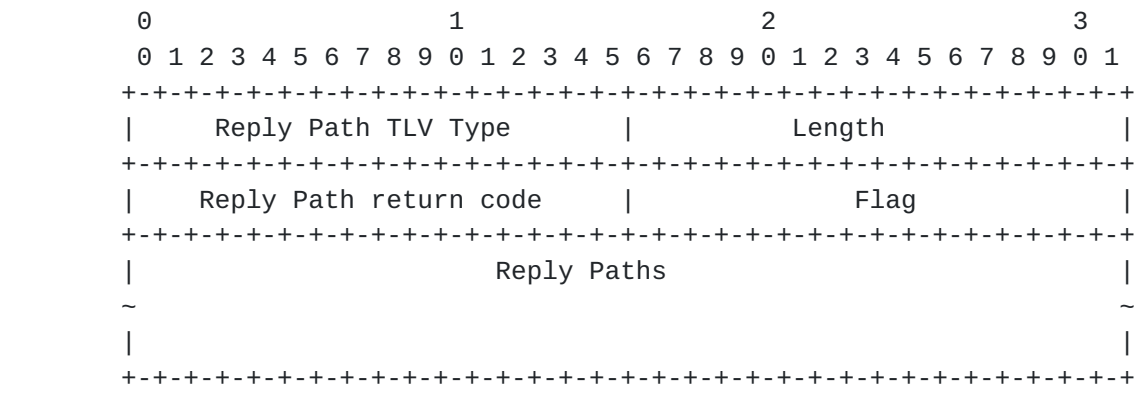

Figure 1 Reply Path TLV

Reply Path TLV Type field is 2 octets in length, and the type value is TBD by IANA.

The Length field is 2 octets in length. It defines the length in octets of the Reply Path return code, Flag and Reply Paths fields.

Reply Path return code is 2 octets in length. It is defined for the egress LSR of the forward LSP to report the results of Reply Path TLV processing and return path selection. When sending echo request,

these codes MUST be set to zero. Reply Path return code only used when sending echo reply, and it MUST be ignored when processing echo request message. This document defines the following Reply Path return codes:

Value	Meaning
0x0000	No return code
0x0001	Malformed Reply Path TLV was received
0x0002	One or more of the sub-TLVs in Reply Path TLV was not understood
0x0003	The echo reply was sent successfully using the specified Reply Path
0x0004	The specified Reply Path was not found, the echo reply was sent via other LSP
0x0005	The specified Reply Path was not found, the echo reply was sent via IP path
0x0006	The Reply mode in echo request was not set to 5(Reply via Specified Path) although Reply Path TLV exists
0x0007	Reply Path TLV was missing in echo request
0x0008-0xfffb	Not allocated, allocated via Standard Action
0xfffc-0xffff	Experimental Use

Flag field is also 2 octets in length, it is used to notify the egress how to process the Reply Paths field when performing return path selection. The Flag field is a bit vector and has following format:

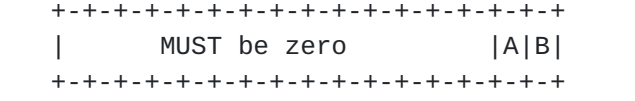

A (Alternative path): the egress LSR MUST select a non-default path as the return path. This is very useful when reverse default path problems are suspected which can be confirmed when the echo reply is forced to follow a non-default return path. Here, the default path refers to the path that the egress LSR will use to send the echo reply when the return path is not explicitly specified as defined in this document. If A bit is set, there is no need to carry any specific reply path sub-TLVs, and when received, the sub-TLVs SHOULD be ignored.

B (Bidirectional): the return path is required to follow the reverse direction of the tested bidirectional LSP. If B bit is set, there is no need to carry any specific reply path sub-TLVs, and when received, the sub-TLVs SHOULD be ignored.

The A bit and B bit set MUST NOT both be set, otherwise, an echo reply with the RP return code set to "Malformed RP TLV was received" SHOULD be returned.

The Reply Paths field is variable in length, not more than one sub-TLV MUST be carried, which describes the specified path that the echo reply message is required to follow. When the Reply Mode field is set to "Reply via Specified Path" in an LSP echo request message, the Reply Path TLV MUST be present. The Reply Path TLV SHALL only be used in the reply mode defined in this document (Reply via Specified Path).

3.3. Reply Path sub-TLVs

Each of the FEC sub-TLVs (include existing and future defined) for the Target FEC Stack TLV[RFC4379] is applicable to be a sub-TLV for inclusion in the Reply Path TLV for expressing a specific return path. For these shared sub-TLVs, they share the same registry with the Target FEC Stack TLV for the range of 0-31743 and 32768-64511.

In addition, this document defines three new sub-TLVs: IPv4 RSVP Tunnel sub-TLV, IPv6 RSVP Tunnel sub-TLV and Static Tunnel sub-TLV. These sub-TLVs are only designed for Reply Path TLV, hence this document calls them dedicated sub-TLVs to Reply Path TLV. For these dedicated sub-TLVs, this document will create a new registry (Section 6.1), the sub-TLV type MUST be allocated from the new registry. Detailed definition is in the following sections.

In [RFC4379], the range of 31744-32767 and 64512-65535 for sub-TLVs is specified for Vendor Private Use, and MUST NOT be allocated. This document changes that rule to make it not applicable to Reply Path TLV and redefines the rule as in Section 6.2 . If an implementation recognizes any specific Vendor Private types as defined in [RFC4379], and uses the sub-TLV type specified in this document, care must be taken to ensure that the implementation does not confuse the two usages.

With the Return Path TLV flags and the sub-TLVs defined for the Target FEC Stack TLV and in this document, it could provide following options for return paths specifying:

1. Specify a particular LSP as return path

   - use those sub-TLVs defined for the Target FEC Stack TLV

2. Specify a more generic tunnel FEC as return path

- use the IPv4/IPv6 RSVP and Static Tunnel sub-TLVs defined in Section 3.3.1, Section 3.3.2 and Section 3.3.3 of this document

3. Specify the reverse path of the bidirectional LSP as return path

   - use B bit defined in Section 3.2 of this document.

4. Force return path to non-default path

   - use A bit defined in Section 3.2 of this document.

3.3.1. IPv4 RSVP Tunnel sub-TLV

The IPv4 RSVP Tunnel sub-TLV is used in the Reply Path TLV to allow the operator to specify a more generic tunnel FEC other than a particular LSP as the return path. According to the bits set in the Flag field, the egress LSR will then choose an LSP from the specified Tunnel as the return path. The format of IPv4 RSVP Tunnel sub-TLV is as follows:

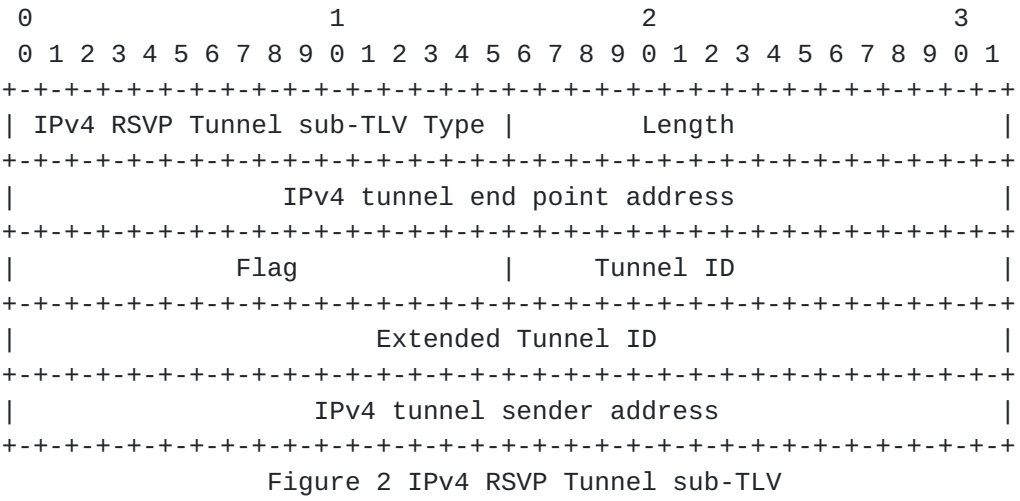

Figure 2 IPv4 RSVP Tunnel sub-TLV

The IPv4 RSVP Tunnel sub-TLV is derived from the RSVP IPv4 FEC TLV that is defined in Section 3.2.3 [RFC4379]. All fields have the same semantics as defined in [RFC4379] except that the LSP-ID field is omitted and a new Flag field is defined.

The IPv4 RSVP Tunnel sub-TLV Type field is 2 octets in length, and the recommended type value is TBD.

The Flag field is 2 octets in length, it is used to notify the egress LSR how to choose the return path. The Flag field is a bit vector and has following format:

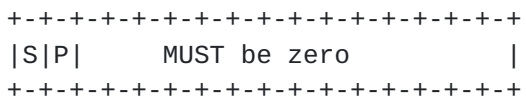

P (Primary): the return path MUST be chosen from the LSPs that belong to the specified Tunnel and the LSP MUST be the primary LSP.

S (Secondary): the return path MUST be chosen from the LSPs that belong to the specified Tunnel and the LSP MUST be the secondary LSP.

P bit and S bit MUST NOT both be set, otherwise, an echo reply with the RP return code set to "Malformed RP TLV was received" SHOULD be returned. If P bit and S bit are both not set, the return path could be any one of the LSPs from the same Tunnel.

3.3.2. IPv6 RSVP Tunnel sub-TLV

The IPv6 RSVP Tunnel sub-TLV is used in the Reply Path TLV to allow the operator to specify a more generic tunnel FEC other than a particular LSP as the return path. According to the bits set in the Flag field, the egress LSR will then choose an LSP from the specified Tunnel as the return path. The format of IPv6 RSVP Tunnel sub-TLV is as follows:

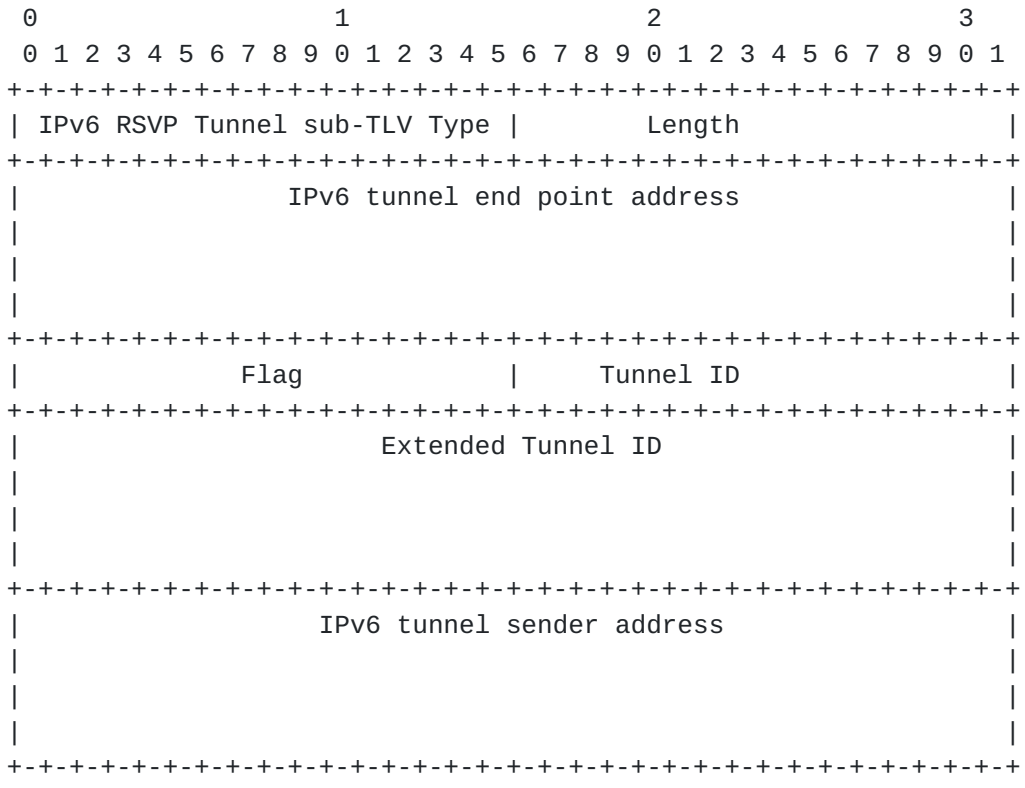

Figure 3 IPv6 RSVP Tunnel sub-TLV

The IPv6 RSVP Tunnel sub-TLV is derived from RSVP IPv6 FEC TLV that is defined in Section 3.2.4 of [RFC4379].All fields have the same semantics as defined in [RFC4379] except that the LSP-ID field is omitted and a new Flag field is defined.

The IPv6 RSVP Tunnel sub-TLV Type field is 2 octets in length, and the type value is TBD.

The Flag field is 2 octets in length and is identical to that described in Section 3.3.1.

3.3.3. Static Tunnel sub-TLV

The Static Tunnel sub-TLV is used in the Reply Path TLV to allow the operator to specify a more generic tunnel FEC other than a particular LSP as the return path. According to the bits set in the Flag field, the egress LSR will then choose an LSP from the specified Tunnel as the return path. The format of Static RSVP Tunnel sub-TLV is as follows. The value fields are taken from the definitions in [RFC6370].

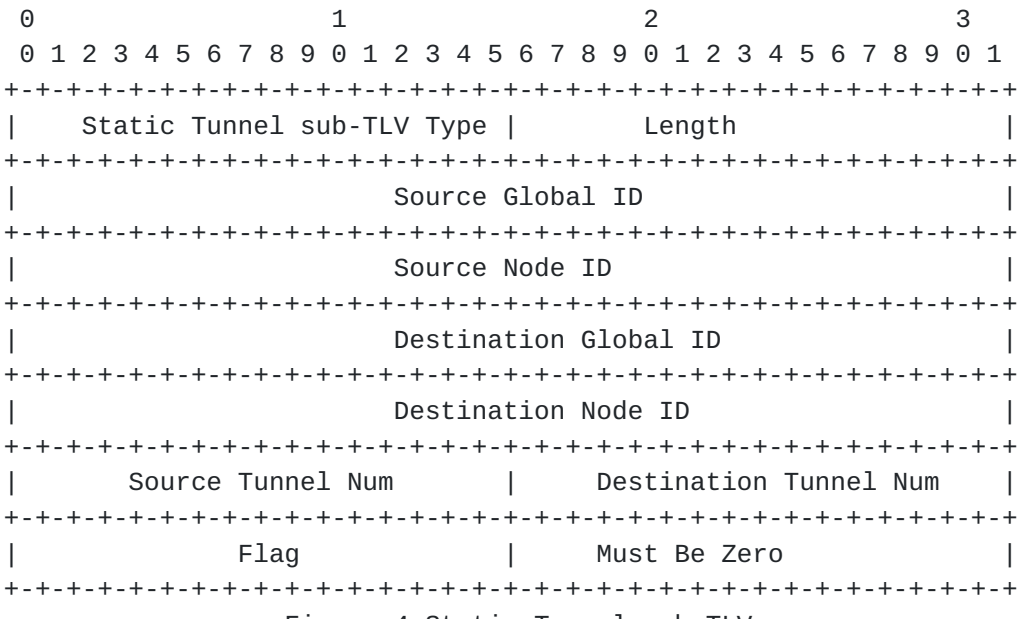

Figure 4 Static Tunnel sub-TLV

The Flag field is 2 octets in length and is identical to that described in Section 3.3.1.

The sub-TLV type value is TBD.

3.4. Reply TC TLV

Reply TOS Byte TLV [RFC4379] is used by the originator of the echo request to request that an echo reply be sent with the IP header TOS byte set to the value specified in the TLV. Similarly, in this document, a new TLV: Reply TC TLV is defined and MAY be used by the originator of the echo request to request that an echo reply be sent with the TC bits of the return path LSP set to the value specified in this TLV. The Reply TC TLV is not limited to the reply mode specified in this document (Reply via Specified Path) but may be used in all the other reply modes as well. The format of Reply TC TLV is as follows:

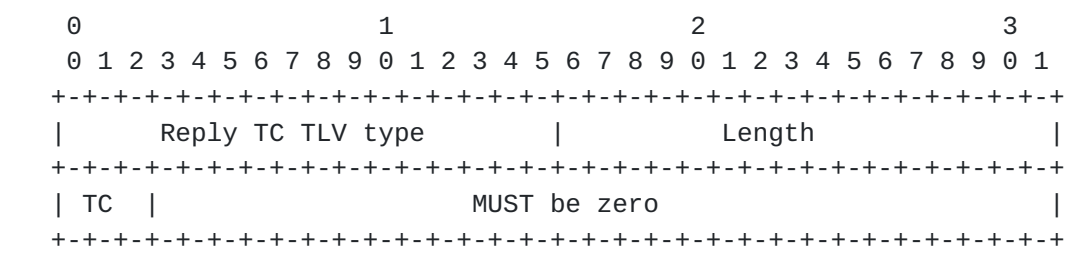

The Reply TC TLV Type field is 2 octets in length, and the type value is TBD.

The Length field is 2 octets in length, the value of length field is fixed 4 octets.

4. Theory of Operation

The procedures defined in this document currently only apply to "ping" mode. The "traceroute" mode is out of scope for this document.

In [RFC4379], the echo reply is used to report the LSP checking result to the LSP Ping initiator. This document defines a new reply mode and a new TLV (Reply Path TLV) that enable the LSP ping initiator to specify or constrain the return path of the echo reply. Similarly the behavior of echo reply is extended to detect the requested return path by looking at a specified path FEC TLV. This enables LSP Ping to detect failures in both directions of a path with a single operation, this of course cuts in half the operational steps required to verify the end to end bidirectional connectivity and integrity of an LSP.

When the echo reply message is intended to test the return MPLS LSP path(when the A bit is not set in the previous received echo request message), the destination IP address of the echo reply message MUST

never be used in a forwarding decision. To avoid this possibility the destination IP address of the echo reply message that is transmitted along the specified return path MUST be set to numbers from the range 127/8 for IPv4 or 0:0:0:0:0:FFFF:127/104 for IPv6, and the IP TTL MUST be set 1, and the TTL in the outermost label MUST be set to 255. Of course when the echo reply message is not intended for testing the specified return path (when the A bit is set in the previous received echo request message) , the procedures defined in [RFC4379] (the destination IP address is copied from the source IP address) apply unchanged.

4.1. Sending an Echo Request

When sending an echo request, in addition to the rules and procedures defined in Section 4.3 of [RFC4379], the reply mode of the echo request MUST be set to "Reply via Specified Path", and a Reply Path TLV MUST be carried in the echo request message correspondingly. The Reply Path TLV includes one or several reply path sub-TLV(s) to identify the return path(s) the egress LSR should use for its reply.

For a bidirectional LSP, since the ingress LSR and egress LSR of a bidirectional LSP are aware of the relationship between the forward and backward direction LSPs, only the B bit SHOULD be set in the Reply Path TLV. If the operator wants the echo reply to be sent along a different path other than the reverse direction of the bidirectional LSP, the "A" bit SHOULD be set or another FEC sub-TLV SHOULD be carried in the Reply Path TLV instead, and the B bit MUST be clear.

In some cases, operators may want to treat two unidirectional LSPs (one for each direction) as a pair. There may not be any binding relationship between the two LSPs. Using the mechanism defined in this document, operators can run LSP Ping one time from one end to complete the failure detection on both unidirectional LSPs. To accomplish this, the echo request message MUST carry (in the Reply Path TLV) a FEC sub-TLV that belongs to the backward LSP.

4.2. Receiving an Echo Request

"Ping" mode processing as defined in Section 4.4 of [RFC4379] applies in this document. In addition, when an echo request is received, if the egress LSR does not know the reply mode defined in this document, an echo reply with the return code set to "Malformed echo request" and the Subcode set to zero will be send back to the ingress LSR according to the rules of [RFC4379]. If the egress LSR knows the reply mode, according to the Reply Path TLV, it SHOULD find and select the desired return path. If there is a matched path, an echo reply with Reply Path TLV that identify the return path SHOULD be

sent back to the ingress LSR, the Reply Path return code SHOULD be set to "The echo reply was sent successfully using the specified return path". If there is no such path, an echo reply with Reply Path TLV SHOULD be sent back to the ingress LSR, the Reply Path return code SHOULD be set to relevant code (defined Section 3.2) for the real situation to reflect the result of Reply Path TLV processing and return path selection. For example, if the specified LSP is not found, the egress then chooses another LSP as the return path to send the echo reply, the Reply Path return code SHOULD be set to "The specified reply path was not found, the echo reply was sent via other LSP", and if the egress chooses an IP path to send the echo reply, the Reply Path return code SHOULD be set to "The specified reply path was not found, the echo reply was sent via IP path". If there is unknown sub-TLV in the received Reply Path TLV, the Reply Path return code SHOULD be set to "One or more of the sub-TLVs in Reply Path TLV was not understood".

If the A bit of the Reply Path TLV in a received echo request message is set, the egress LSR SHOULD send the echo reply along an non-default return path.

IF the B bit of the Reply Path TLV in a received echo request message is set, the egress LSR SHOULD send the echo reply along the reverse direction of the bidirectional LSP.

If the A bit of the Reply Path TLV in a received echo request message is not set(a.k.a a specific return path sub-TLV is carried or the B bit is set), the echo reply is REQUIRED not only to send along the specified path, but to test the selected return path as well (by carrying the FEC stack information of the return path).

In addition, the FEC validate results of the forward path LSP SHOULD NOT affect the egress LSR continue to test return path LSP.

4.3. Sending an Echo Reply

As described in [RFC4379], the echo reply message is a UDP packet, and it MUST be sent only in response to an MPLS echo request. The source IP address is a valid IP address of the replier, the source UDP port is the well-know UDP port for LSP ping.

When the echo reply is intended to test the return path (the A is not set in the previous received echo request), the destination IP address of the echo reply message MUST never be used in a forwarding decision. To avoid this problem, the IP destination address of the echo reply message that is transmitted along the specified return path MUST be set to numbers from the range 127/8 for IPv4 or 0:0:0:0:0:FFFF:127/104 for IPv6, and the IP TTL MUST be set to 1, the TTL in

the outermost label MUST be set to 255. Otherwise, the same as defined in [RFC4379], the destination IP address and UDP port are copied from the source IP address and source UDP port of the echo request.

When sending the echo reply, a Reply Path TLV that identifies the return path MUST be carried, the Reply Path return code SHOULD be set to relevant code that reflects results about how the egress processes the Reply Path TLV in a previous received echo request message and return path selection. By carrying the Reply Path TLV in an echo reply, it gives the Ingress LSR enough information about the reverse direction of the tested path to verify the consistency of the data plane against control plane. Thus a single LSP Ping could achieve both directions of a path test. If the return path is pure IP path, no sub-TLVs are carried in the Reply Path TLV.

4.4. Receiving an Echo Reply

The rules and process defined in Section 4.6 of [RFC4379] apply here. When an echo reply is received, if the reply mode is "Reply via Specified Path" and the Reply Path return code is "The echo reply was sent successfully using the specified return path", and if the A bit is not set. The ingress LSR MUST perform FEC validation (based on the FEC stack information of the return path carried in the Reply Path TLV) as an egress LSR does when receiving an echo request, the FEC validation process (relevant to "ping" mode) defined in Section 4.4.1 of [RFC4379] applies here.

When an echo reply is received with return code set to "Malformed echo request received" and the Subcode set to zero. It is possible that the egress LSR may not know the "Reply via Specified Path" reply mode, the operator may choose to re-perform another LSP Ping by using one of the four reply modes defined [RFC4379].

On receipt of an echo reply with Reply Path return code in the Reply Path TLV set to "The specified reply path was not found, ...", it means that the egress LSR could not find a matched return path as specified. Operators may choose to specify another LSP as the return path or use other methods to detect the path further.

5. Security Considerations

Security considerations discussed in [RFC4379] apply to this document. In addition to that, in order to prevent using the extension defined in this document for "proxying" any possible attacks, the return path LSP MUST have destination to the same node where the forward path is from.

6. IANA Considerations

IANA has a temporary allocation for a TLV from the "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters - TLVs" registry, "TLVs and sub-TLVs" sub-registry - Type 21 (Reply Path TLV). For this TLV the standards action sub-TLVs (the range of 0-31743 and 32768-64511) shall be blocked from being allocated. IANA is also requested to assign one new TLV from the "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters - TLVs" registry, "TLVs and sub-TLVs" sub-registry.

6.1. Temporary assigned TLV and New TLV

The IANA is requested to assign the temporary assigned Reply path TLV and also assign one new TLV from the "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters - TLVs" registry, "TLVs and sub-TLVs" sub-registry.

Note: IANA have made an early allocation of the value 21 for Reply Path TLV.

Value	Meaning	Reference
21	Reply Path TLV	this document (sect 3.2)
TBD	Reply TC TLV	this document (sect 3.4)

6.2. Sub-TLVs

The sub-TLV range of Reply Path TLV are partitioned as following:

0-31743 - Reserved, and MUST NOT be allocated.

31744-32767 - Allocated via Standards Action.

32768-64511 - Reserved, and MUST NOT be allocated.

64512-65531 - Allocated via Standards Action.

65531-65535 - Experimental Use, and MUST NOT be allocated.

6.2.1. Dedicated Sub-TLVs to Reply Path TLV

IANA is also requested to assign three new sub-TLV types from "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters - TLVs" registry, "TLVs and sub-TLVs" sub-registry for the Reply Path TLV (Type 21) - from the Standards Action range.

Sub-type	Value Field	Reference
TBD	IPv4 RSVP Tunnel	this document (sect 3.3.1)
TBD	IPv6 RSVP Tunnel	this document (sect 3.3.2)
TBD	Static Tunnel	this document (sect 3.3.3)

Previously temporary allocated sub-TLVs fall within the blocked range an should NOT be allocated.

6.3. New Reply Mode

IANA is now requested to assign the previously assigned a new reply mode code point (5 - Reply via specified path) from the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry, the "Reply Mode" sub-registry on a permanent basis.

Value	Meaning	Reference
5	Reply via Specified Path	this document (sect 3.1)

6.4. Reply Path Return Code

IANA is requested to create a new registry for Reply Path return code.

This document (Section 3.2) defines the following return codes:

Value	Meaning
0x0000	No return code
0x0001	Malformed Reply Path TLV was received
0x0002	One or more of the sub-TLVs in Reply Path TLV was not understood
0x0003	The echo reply was sent successfully using the specified Reply Path
0x0004	The specified Reply Path was not found, the echo reply was sent via other LSP
0x0005	The specified Reply Path was not found, the echo reply was sent via IP path
0x0006	The Reply mode in echo request was not set to 5(Reply via Specified Path) although Reply Path TLV exists
0x0007	Reply Path TLV was missing in echo request
0x0008-0xfffb	Not allocated, allocated via Standard Action
0xfffc-0xffff	Experimental Use

The range of 0x0008-0xfffb is not allocated and reserved for future

extensions and is allocated via Standard Action, the range of 0xfffc-0xffff is for Experimental Use.

7. Contributors

The following individuals also contributed to this document:

Ehud Doron
Orckit-Corrigent
EMail: ehudd@orckit.com

Ronen Solomon
Orckit-Corrigent
EMail: RonenS@orckit.com

Ville Hallivuori
Tellabs
Sinimaentie 6 C
FI-02630 Espoo, Finland
EMail: ville.hallivuori@tellabs.com

Xinchun Guo
EMail: guoxinchun@huawei.com

8. Acknowledgements

The authors would like to thank Adrian Farrel, Peter Ashwood-Smith, Sriganesh Kini, Gregory Mirsky, Eric Gray, Loa Andersson and Tom Petch for their review, suggestion and comments to this document.

9. References

9.1. 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.

[RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, September 2011.

9.2. Informative References

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, February 2009.

[RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

Authors' Addresses

Mach(Guoyi) Chen
Huawei Technologies Co., Ltd
Q14 Huawei Campus, No. 156 Beiqing Road, Hai-dian District
Beijing 100095
China

Email: mach@huawei.com

Wei Cao
Huawei Technologies Co., Ltd
Q14 Huawei Campus, No. 156 Beiqing Road, Hai-dian District
Beijing 100095
China

Email: wayne.caowei@huawei.com

So Ning
Tata Communications

Email: ning.so@tatacommunications.com

Frederic Jounay
Orange CH
4 rue caudray 1020 Renens
Switzerland

Email: frederic.jounay@orange.ch

Simon Delord
Alcatel-Lucent
Building 3, 388 Ningqiao Road, Jinqiao, Pudong
Shanghai 201206
China

Email: simon.delord@alcatel-lucent.com