

Network Working Group
Internet Draft
Category: Standards Track
Expiration Date: September 2007

George Swallow
Cisco Systems, Inc.

Vanson Lim
Cisco Systems, Inc.

March 2007

Proxy LSP Ping

[draft-swallow-mpls-remote-lsp-ping-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Abstract

This document defines a means of remotely initiating Multiprocal Label Switched Protocol Pings on Label Switched Paths. A proxy ping request is sent to any Label Switching Routers along a Label Switched Path. The primary motivations for this facility are first to limit the number of messages and related processing when using LSP Ping in large Point-to-Multipoint LSPs, and second to enable leaf to root tracing.

Contents

1	Introduction	3
1.1	Conventions	3
2	Proxy Ping Overview	4
3	Remote Echo / Reply Pprocedures	5
3.1	Procedures for the initiator	5
3.2	Procedures for the proxy LSR	6
3.2.1	Sending an MPLS proxy ping reply	7
3.2.2	Sending the MPLS echo requests	7
4	Proxy Ping Request / Reply Messages	8
4.1	Proxy Ping Request / Reply Message formats	9
4.2	Proxy Ping Request Message contents	10
4.3	Proxy Ping Reply Message Contents	10
5	Object formats	11
5.1	Proxy Echo Parameters Object	11
5.2	Previous Hop Object	13
6	Security Considerations	13
7	IANA Considerations	13
7.1	Message and Object Type Assignments	13
7.2	Return Code Assignments	14
8	Acknowledgments	14
9	References	14
9.1	Normative Referenes	14
9.2	Informative References	14
10	Authors' Addresses	15

1. Introduction

It is anticipated that very large Point-to-Multipoint (P2MP) Label Switched Paths (LSPs) will exist. Further it is anticipated that many of the applications for P2MP tunnels will require OAM that is both rigorous and scalable.

Suppose one wishes to trace a P2MP LSP to localize a fault which is affecting one egress or a set of egresses. Suppose one follows the normal procedure for tracing - namely repeatedly pinging from the root, incrementing the TTL by one after each three or so pings. Such a procedure has the potential for producing a large amount of processing at the P2MP-LSP midpoints and egresses. It also could produce an unwieldy number of replies back to the root.

One alternative would be to begin sending pings from points at or near the affected egress(es) and working backwards toward the root. The TTL could be held constant as say two, limiting the the number of responses to the number of next-next-hops of the point where the ping was initiated.

The above procedure does require that the root know the previous-hop node to the one which was pinged on the prior iteration. This information is readily available in [[P2MP-TE](#)]. This document provides a means for obtaining this information for [[mLDP](#)] as well as defining a means for remotely causing an MPLS echo request message to be sent down a Label Switched Path (LSP) or part of an LSP.

While the motivaton for this document came from multicast scaling concerns, its applicability may be wider. However other uses of this facility are beyond the scope of this document. Further the discussion is cauched in terms of multipoint LSPs.

1.1. Conventions

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](#) [[KeyWords](#)].

The term "Must Be Zero" (MBZ) is used in object descriptions for reserved fields. These fields MUST be set to zero when sent and ignored on receipt.

Based on context the terms leaf and egress are used interchangeably. Egress is used where consistency with [[RFC4379](#)] was deemed appropriate. Receiver is used in the context of receiving protocol messages.

[Note (to be removed after assignments occur): <tba> = to be assigned by IANA]

2. Proxy Ping Overview

Two new LSP Ping messages are defined for remote pinging, the MPLS proxy ping request and the MPLS proxy ping reply.

A remote ping operation on a P2MP LSP involves at least three LSRs; in some scenarios none of these are the ingress (root) or an egress (leaf) of the LSP.

We refer to these nodes with the following terms:

Initiator - the node which initiates the ping operation by sending an MPLS proxy ping request message

Proxy LSR - the node which is the destination of the MPLS proxy request message and potential initiator of the MPLS echo request

Receiver(s) - the receivers of the MPLS echo request messages

The initiator formats an MPLS proxy ping request message and sends it to the proxy LSR, a node it believes to be on the path of the LSP. This message specifies the MPLS echo request to be sent inband of the LSP. It may also request the proxy LSR to acknowledge the receipt of the proxy ping request message and/or respond with the address of the previous hop, i.e. the LSR upstream of it on this LSP.

The proxy LSR validates that it has a label mapping for the specified FEC and that it is authorized to send the specified MPLS echo request on behalf of the initiator. Depending on the Reply Mode carried in the header of the proxy ping request message and the above results an MPLS remote echo reply message might be sent back to the initiator. This message may also communicate the address of the previous hop.

If the proxy LSR has a label mapping for the FEC and all authorization check have passed, the proxy LSR formats an MPLS echo request. If the source address of the IP packet is not the initiator, it includes a ReplyTo object containing the initiator's address. It then sends it inband of the LSP.

The receivers process the MPLS echo request as normal, sending their MPLS echo replies back to the initiator.

3. Remote Echo / Reply Procedures

3.1. Procedures for the initiator

The initiator creates an MPLS proxy ping request message.

The message MUST contain a Target FEC Stack that describes the FEC being tested.

[Note for the current version of the ID, the FEC stack is limited to a single FEC as we have not yet fully considered the operational and security impacts of permitting more FECs]

The message MUST contain a Proxy Echo Parameters object. The address type is set to either IPv4 or IPv6. The Destination IP Address is set to the value to be used in the MPLS echo request packet. If the Address Type is IPv4, an address from the range 127/8. If the Address Type is IPv6, an address from the range 0:0:0:0:0:FFFF:127/104. By default the source address will be set to an address of the proxy LSR.

The Reply mode and Global Flags of the Proxy Echo Parameters object are set to the values to be used in the MPLS echo request message header. The Source UDP Port is set to the value to be used in the MPLS echo request packet. The TTL is set to the value to be used in the outgoing MPLS label stack. See [section 5.2.2.2](#) for further details.

Flags MAY be set to request the previous hop address and/or a downstream mapping object from the proxy LSR.

A list of Next Hop IP Addresses MAY be included to limit the next hops towards which the MPLS echo request message will be sent.

Any of following objects MAY be included; these objects will be copied into the MPLS echo request messages:

- Pad
- Vendor Enterprise Number
- Reply TOS Byte
- P2MP Egress Identifier [[McstPing](#)]
- Echo Jitter TLV [[McstPing](#)]
- Vendor Private TLVs

Downstream Mapping objects MAY be included. These objects will be matched to the next hop address for inclusion in those particular MPLS echo request messages.

The message is then encapsulated in a UDP packet. The source UDP port is chosen by the sender; the destination UDP port is set to 3503. The IP header is set as follows: the source IP address is a routable address of the sender; the destination IP address is a routable address of the midpoint. The packet is then sent with the IP TTL is set to 255.

3.2. Procedures for the proxy LSR

A proxy LSR that receives an MPLS proxy ping request message, parses the packet to ensure that it is a well-formed packet. It checks that the TLVs that are not marked "Ignore" are understood. If not, it sets the Return Code set to "Malformed echo request received" or "TLV not understood" (as appropriate), and the Subcode set to zero. If the Reply Mode of the message header is not 0, an MPLS proxy ping reply message SHOULD be sent as described below. In the latter case, the misunderstood TLVs (only) are included in an Errored TLVs object.

The header fields Sender's Handle and Sequence Number are not examined, but are saved to be included in the MPLS proxy ping reply and MPLS echo request messages.

The proxy LSR validates that it has a label mapping for the specified FEC, it then determines if it is an egress, transit or bud node and sets the Return Code as appropriate.

The proxy LSR then determines if it is authorized to send the specified MPLS echo request on behalf of the initiator. An LSR MUST be capable of filtering addresses to validate initiators. Other filters on FECs or MPLS echo request contents MAY be applied. If a filter has been invoked (i.e. configured) and an address does not pass the filter, then an MPLS echo request message MUST NOT be sent, and the event SHOULD be logged. An MPLS proxy ping reply message may be sent with a Return Code of <tba>, "Remote Ping not authorized".

If the "Request for Previous Hop" flag is set, a PHOP Address Object is formatted for inclusion in the MPLS proxy ping reply. If the previous HOP is unknown or ambiguous the Address Type is set to "No Address Supplied".

If there is a list of Next Hop addresses in the Proxy Echo Parameters object, each address is examined to determine if it is a next hop for this FEC. If any are not, those addresses are deleted from the list. The updated Proxy Echo Parameters object is included in the MPLS proxy ping reply.

If the "Request for Downstream Mapping" flag is set the LSR formats a

Downstream Mapping object for each interface that the MPLS echo request will be sent out.

If the Reply Mode of the message header is not 0 or 5, an MPLS remote echo reply message SHOULD be sent as described below.

3.2.1. Sending an MPLS proxy ping reply

The Reply mode, Sender's Handle and Sequence Number fields are copied from the proxy ping request message. Various objects are included as specified above. The message is encapsulated in a UDP packet. The source IP address is a routable address of the proxy LSR; the source port is the well-known UDP port for LSP ping. The destination IP address and UDP port are copied from the source IP address and UDP port of the echo request. The IP TTL is set to 255.

3.2.2. Sending the MPLS echo requests

A base MPLS echo request is formed as described in the next section. The section below that describes how the base MPLS echo request is sent on each interface.

3.2.2.1. Forming the base MPLS echo request

A Next_Hop_List is created as follows. If Next Hop addresses were included in the received Proxy Parameters object, the Next_Hop_List is copied from the Proxy Echo Parameters object as adjusted above. Otherwise, the list is set to all the next hops to which the FEC would be forwarded.

The proxy LSR then formats an MPLS echo request message. The Global Flags and Reply Mode are copied from the Proxy Echo Parameters object. The Return Code and Return Subcode are set to zero.

The Sender's Handle and Sequence Number are copied from the remote echo request message.

The TimeStamp Sent is set to the time-of-day (in seconds and microseconds) that the echo request is sent. The TimeStamp Received is set to zero.

A ReplyTo object (see [[SelfTest](#)]) containing the initiator's address is included.

The following objects are copied from the MPLS proxy ping request

message. Note that of these, only the Target FEC Stack is REQUIRED to appear in the MPLS proxy ping request message.

- Target FEC Stack
- Pad
- Vendor Enterprise Number
- Reply TOS Byte
- P2MP Egress Identifier [[McstPing](#)]
- Echo Jitter TLV [[McstPing](#)]
- Vendor Private TLVs

The message is then encapsulated in a UDP packet. The source UDP port is copied from the Proxy Echo Parameters object. destination ports are copied from the proxy ping request message.

The source IP address is set to a routable address of the proxy LSR. Per usual the TTL of the IP packet is set to 1.

3.2.2.2. Per interface sending procedures

The proxy LSR now iterates through the Next_Hop_List modifying the base MPLS echo request to form the MPLS echo request packet which is then sent on that particular interface.

For each next hop address, the outgoing label stack is determined. The TTL for the label corresponding to the FEC in the FEC stack is set such that the TTL on the wire will be one less than the TTL specified in the proxy ping request message. If any additional labels are pushed onto the stack, their TTLs are set to 255.

If the MPLS proxy ping request message contained Downstream Mapping objects, they are examined. If the Downstream IP Address matches the next hop address that Downstream Mapping object is included in the MPLS echo request.

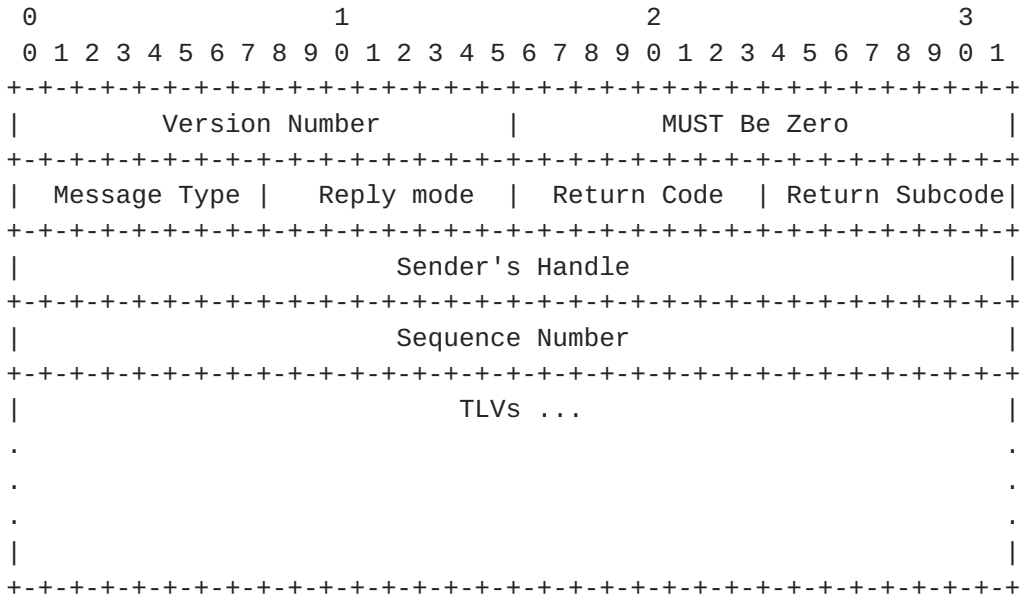
The packet is then transmitted on this interface.

4. Proxy Ping Request / Reply Messages

Two new LSP Ping messages are defined for remote pinging, the MPLS proxy ping request message and the MPLS proxy ping reply.

4.1. Proxy Ping Request / Reply Message formats

Except where noted, the definitions of all fields in the messages are identical to those found in [LSP-PING]. The messages have the following format:



Version Number

The Version Number is currently 1. (Note: the Version Number is to be incremented whenever a change is made that affects the ability of an implementation to correctly parse or process an MPLS echo request/reply. These changes include any syntactic or semantic changes made to any of the fixed fields, or to any TLV or sub-TLV assignment or format that is defined at a certain version number. The Version Number may not need to be changed if an optional TLV or sub-TLV is added.)

Message Type

Type	Message
----	-----
5	MPLS proxy ping request
6	MPLS proxy ping reply

Reply mode

The reply modes are the same as [LSP-PING] with the addition of value 5. For completeness, the full list of reply modes follows:

Value	Meaning
-----	-----
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
5	Reply via an IPv4/IPv6 UDP packet only if the proxy request is not fulfilled

4.2. Proxy Ping Request Message contents

The MPLS proxy ping request message MAY contain the following objects:

Type	Object
-----	-----
1	Target FEC Stack
2	Downstream Mapping
3	Pad
5	Vendor Enterprise Number
10	Reply TOS Byte
tba	Proxy Echo Parameters
tba	PHOP Address
tba	P2MP Egress Identifier [McstPing]
tba	Echo Jitter TLV [McstPing]
	Vendor Private TLVs

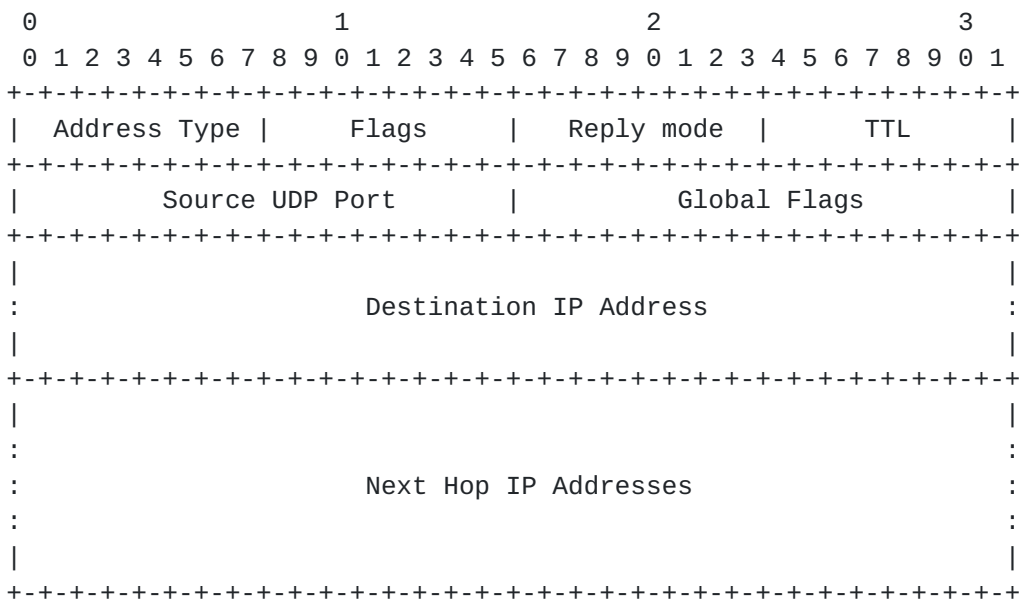
4.3. Proxy Ping Reply Message Contents

The MPLS proxy ping reply message MAY contain the following objects:

Type	Object
----	-----
1	Target FEC Stack
2	Downstream Mapping
5	Vendor Enterprise Number
9	Errored TLVs
tba	Proxy Echo Parameters
tba	PHOP Address
	Vendor Private objects

5. Object formats

5.1. Proxy Echo Parameters Object



Address Type

The type of address found in the in the Destination IP Address and Next Hop IP Addresses fields. The type codes appear in the table below:

Address Family	Type
IPv4 Numbered	1
IPv6 Numbered	3

Flags

Request for Previous Hop 0x01

When set this requests that the proxy LSR supply the PHOP address in the MPLS proxy ping reply message

Request for Downstream Mapping 0x02

When set this requests that the proxy LSR supply a Downstream Mapping object in the MPLS proxy ping reply message

Reply mode

The reply mode to be sent in the MPLS Echo Request message; the values are as specified in [[RFC4379](#)]

TTL

The TTL to be used in the label corresponding to the FEC in the MPLS Echo Request packet

Source UDP Port

The source UDP port to be sent in the MPLS Echo Request packet

Global Flags

The Global Flags to be sent in the MPLS Echo Request message

Destination IP Address

If the Address Type is IPv4, an address from the range 127/8;
If the Address Type is IPv6, an address from the range
0:0:0:0:0:FFFF:127/104

Next Hop IP Addresses

A list of next hop address that the echo request message is to be sent towards

5.2. Previous Hop Object

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Address Type |                               MUST be Zero                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|
:                               Previous Hop IP Address                               :
|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
    
```

Address Type

A type code as specified in the table below:

Type	Type of Address
0	No Address Supplied
1	IPv4
3	IPv6

6. Security Considerations

[To be written]

7. IANA Considerations

[Not complete]

7.1. Message and Object Type Assignments

This document makes the following codepoint assignments (pending IANA action):

Registry	Codepoint	Purpose
LSP Ping Message Type	tba	MPLS proxy ping request message
	tba	MPLS proxy ping reply
LSP Ping Object Type	tba	Proxy Echo Parameters
	tba	PHOP Address

7.2. Return Code Assignments

Value	Meaning
tba	Remote Ping not authorized
tba	Failed Next Hops

8. Acknowledgments

9. References

9.1. Normative References

- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), February 2006.
- [SelfTest] Swallow, G. et al., "LSR Self Test", [draft-ietf-mpls-lsr-self-test-06.txt](#), October 2005.
- [KeyWords] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [McstPing] Farrel, A. et al, "Detecting Data Plane Failures in Point-to-Multipoint MPLS Traffic Engineering - Extensions to LSP Ping", [draft-ietf-mpls-p2mp-lsp-ping-02.txt](#), September 2006.

9.2. Informative References

- [P2MP-TE] Aggarwal, R., et al., "Extensions to RSVP-TE for Point-to-Multipoint TE LSPs", [draft-ietf-mpls-rsvp-te-p2mp-07.txt](#), January 2007.
- [mLDP] Minei, I., et. al., "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths" [draft-ietf-mpls-ldp-p2mp-02.txt](#), October 2006.

10. Authors' Addresses

George Swallow
Cisco Systems, Inc.
1414 Massachusetts Ave
Boxborough, MA 01719

Email: swallow@cisco.com

Vanson Lim
Cisco Systems, Inc.
1414 Massachusetts Ave
Boxborough, MA 01719

Email: vlim@cisco.com

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Full Copyright Notice

Copyright (C) The IETF Trust (2007). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

