

MPLS
Internet-Draft
Intended status: Standards Track
Expires: October 9, 2014

S. Bryant
S. Sivabalan
S. Soni
Cisco Systems
April 7, 2014

**MPLS Performance Measurement UDP Return Path
draft-bryant-mpls-oam-udp-return-00**

Abstract

This document specifies an extension to the protocol for making performance measurements of MPLS LSPs that is defined in [RFC6374](#). It specifies the procedure used for sending and processing MPLS performance management out-of-band responses for delay and loss measurements over an IP/UDP return path.

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 October 9, 2014.

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

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

1. Introduction

[RFC6374] does not define how an MPLS performance measurement (PM) out-of-band response delivered over IP will be transmitted to the Querier.

In a highly scaled system some PM sessions may be off-loaded to a specific node within a the distributed system that comprises the LSR as a whole. In such systems the response may arrive via any interface in the LSR and need to internally forwarded to the processor tasked with handling the particular PM measurement. Currently the MPLS PM protocol does not have any mechanism to deliver the PM Response message to particular node within a multi-CPU LSR.

The procedure described in this specification shows how to deliver the response over a dynamic UDP port.

2. 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 [[RFC2119](#)].

3. Solution Overview

This document specifies that, unless configured otherwise, the "Return Address TLV" defined in [[RFC6374](#)] SHALL be used to carry the return address. It also defines "Return UDP Port" TLV which, unless configured otherwise SHALL be used to carry the return UDP port. The Return Address TLV and the Return UDP PORT TLV carried in the MPLS-PM query message are used to specify to the Responder how to return the response message.

The procedures defined in this document may be applied to both unidirectional tunnels and Bidirectional LSPs. In this document, the 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) that 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 the MPLS Transport Profile (MPLS-TP).

3.1. Return UDP Port TLV

The format of the Return UDP TLV is as follows:

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| UDP TLV Type |          Length          |      UDP Dest Port      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The UDP TLV Type has a value of <TBD>.

4. Theory of Operation

This document defines the "Return UDP Port TLV" and uses "Return Address TLV" that enables the MPLS-PM Querier to specify the return path for the MPLS-PM reply using IP/UDP encapsulation.

When the MPLS-PM Response is requested out-of-band by setting Control Code of the MPLS-PM Query to "Out-of-band Response Requested", the responder SHOULD send the response back to Querier on the specified destination UDP port at the specified destination IP address as received in the "Return UDP Port TLV and "Return Address TLV" respectively.

If either the "Return Address TLV" or "Return UDP port TLV" is not present in Query Packet and MPLS PM Response is requested out-of-band, the Query message MUST NOT be processed further. If received over a bidirectional LSP, the control code of the Response packet MUST be set to "Invalid Message" and a Response SHOULD be sent over the reverse LSP. The receipt of such a mal-formed request SHOULD be notified to the operator through the management system, taking the normal precautions with respect to the prevention of overload of the error reporting system.

4.1. Sending an MPLS-PM Query

When sending an MPLS PM Query packet, in addition to the rules and procedures defined in [RFC6374]; the Control Code of the MPLS-PM Query MUST be set to "Out-of-band Response Requested", and a "Return UDP Port TLV" along with "Return Address TLV" MUST be carried in the MPLS-PM Query message.

Since the Querier uses the UDP port to de-multiplex response for different measurement type, there SHOULD be a different UDP port for each measurement type (Delay, loss and delay-loss combined).

Implementation MAY use multiple UDP ports for same measurement type to direct the response to the correct management process in the LSR.

4.2. Receiving an MPLS PM Query Request

The processing of MPLS-PM query messages as defined in [[RFC6374](#)] applies in this document. In addition, when an MPLS PM Query request is received, with the Control Code of the MPLS-PM Query set to "Out-of-band Response Requested" with a Return address TLV and Return UDP TLV is present, then the Responder SHOULD use that IP address and UDP port to send MPLS-PM response back to Querier.

If an Out-of-band response is requested and either the Return Address TLV or the Return UDP port TLV is missing, the Query SHOULD be dropped in the case of unidirectional LSP. If either of these TLVs is missing on a bidirectional LSP, the control code of Response packet should set to "Invalid Message" and the response SHOULD be sent over the reverse LSP. In either case the receipt of such a malformed request SHOULD be notified to the operator through the management system, taking the normal precautions with respect to the prevention of overload of the error reporting system.

4.3. Sending an MPLS-PM Response

As specified in [[RFC6374](#)] the MPLS PM Response packet can be sent over either the reverse MPLS LSP for a bidirectional LSP or over an IP path. It MUST NOT be sent other than in response to an MPLS PM Query Packet.

When the requested return path is an IP forwarding path and this method is in use, the destination IP address and UDP port SHOULD be copied from the Return Address TLV and the Return UDP TLV respectively. The source IP address and the source UDP port of Response packet is left to discretion of the Responder subject to the normal management and security considerations. The packet format for the PM response after the UDP header is as specified in [[RFC6374](#)]. As shown in Figure 1 the Associate Channel Header (ACH) [[RFC5586](#)] is not included. The information provided by the ACH is not needed since the correct binding between the Query and Response messages is achieved through the UDP Port and the Session Identifier contained in the [RFC6374](#) message.


```

+-----+
|  IP Header                                |
.   Source Address = Responders IP Address |
.   Destination Address = Return Address TLV.Address |
.   Protocol = UDP                          .
.                                           .
+-----+
|  UDP Header                              |
.   Source Port = As chosen by Responder   .
.   Destination Port = Return UDP Port TLV.UDP Dest Port .
.                                           .
+-----+
|  Message as specified in RFC6374          |
.                                           .
+-----+

```

Figure 1: Response packet Format

If the return path is IP path, only one-way delay or one-way loss measurement can be carried out. In this case timestamps 3 and 4 MUST be zero as specified in [[RFC6374](#)].

4.4. Receiving an MPLS-PM Response

If the response was received over UDP/IP and an out-of-band response was expected, the Response message SHOULD be directed to the appropriate measurement process as determined by the destination UDP Port, and processed using the corresponding measurement type procedure specified in [[RFC6374](#)].

If the Response was received over UDP/IP and an out-of-band response was not requested, that response should be dropped and the event SHOULD be notified to the operator through the management system, taking the normal precautions with respect to the prevention of overload of the error reporting system.

5. Manageability Considerations

The manageability considerations described in [Section 7 of \[RFC6374\]](#) are applicable to this specification. Additional manageability considerations are noted within the elements of procedure of this document.

Nothing in this document precludes the use of a configured UDP/IP return path in a deployment in which configuration is preferred to signalling. In these circumstances the address and UDP port TLVs MAY be omitted from the MPLS PM messages.

6. Security Considerations

The MPLS PM system is not intended to be deployed on the public Internet. It is intended for deployment in well managed private and service provider networks. The security considerations described in [Section 8 of \[RFC6374\]](#) are applicable to this specification and the reader's attention is drawn to the last two paragraphs. Cryptographic measures may be enhanced by the correct configuration of access control lists and firewalls.

There is no additional exposure of information to pervasive monitoring systems observing LSPs that are being monitored.

7. IANA Considerations

IANA is requested to assign a new Optional TLV type from MPLS Loss/Delay Measurement TLV Object Registry contained within the g-ach-parameters parameters registry set.

Code	Description	Reference
TBD	Return UDP Port	[This]

The TLV 131 is recommended

8. Acknowledgements

We acknowledge the contribution of Joseph Chin and Rakesh Gandhi, both with Cisco Systems.

We thank all who have reviewed this text and provided feedback.

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", [RFC 3945](#), October 2004.
- [RFC5586] Bocci, M., Vigoureux, M., and S. Bryant, "MPLS Generic Associated Channel", [RFC 5586](#), June 2009.
- [RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", [RFC 5654](#), September 2009.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", [RFC 6374](#), September 2011.

Authors' Addresses

Stewart Bryant
Cisco Systems

Email: stbryant@cisco.com

Siva Sivabalan
Cisco Systems

Email: msiva@cisco.com

Sagar Soni
Cisco Systems

Email: sagsoni@cisco.com

