

MPLS
Internet-Draft
Intended status: Standards Track
Expires: October 31, 2015

S. Bryant
S. Sivabalan
S. Soni
Cisco Systems
April 29, 2015

RFC6374 UDP Return Path
draft-ietf-mpls-rfc6374-udp-return-path-03

Abstract

This document specifies the procedure to be used by the Packet Loss and Delay Measurement for MPLS Networks protocol defined in [RFC6374](#) when 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 31, 2015.

Copyright Notice

Copyright (c) 2015 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

This document describes how Packet Loss and Delay Measurement for MPLS Networks protocol (MPLS-PLDM) [[RFC6374](#)] out-of-band responses can be delivered to the Querier using UDP/IP.

The use of UDP may be required to support data path management such as passage through firewalls, or to provide the necessary multiplexing needed in bistatic operation where the querier and the collector are not co-located and the collector is gathering the response information for a number of responders. In a highly scaled system some MPLS-PLDM sessions may be off-loaded to a specific node within the distributed system that comprises the Label Switching Router (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 MPLS-PLDM measurement. Currently the MPLS-PLDM protocol does not have any mechanism to deliver the PLDM Response message to particular node within a multi-CPU LSR.

The procedure described in this specification describes how the queryer requests delivery of the MPLS-PLDM response over IP to a dynamic UDP port. It makes no other changes to the protocol and thus does not affect the case where the response is delivered over a MPLS Associated Channel [[RFC5586](#)].

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, if a UDP Return Object (URO) is present in a MPLS-PLDM Query, the responder MUST use the IP address and UDP port in the URO to reply back to the querier. Multiple UROs MAY be present in a MPLS-PLDM Query indicating that an identical responses SHOULD be sent to each address-port pair. A responder MAY be designed or configured to only transmit a single response, in which case the response MUST be sent using the parameters specified in the first URO in the query packet.

The procedures defined in this document may be applied to both unidirectional 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 to the MPLS Transport Profile (MPLS-TP) [RFC5654], [RFC5921]

3.1. UDP Return Object


NOTE TO RFC Editor please delete the following paragraph before publication.

START DELETE

Note to reviewers - We considered a number of approaches to the design. The first was to use the existing address object and a separate UDP object, but concern in the WG was that there may be more than one collector that required this information, and the combined size of the two objects. The next approach considered by the authors was to create a new object by appending a UDP port to the existing generalized address object. However by noting that UDP is only likely to be sent over IP and that it will be a long time before we design a third major version of IP we can compress the object either by having separate IPv4 and an IPv4 objects, or using the address length as the discriminator. The object design below uses the latter approach. The resultant combined UDP port + address object is thus the same size as the original address object.

END DELETE

The format of the UDP Return Object (URO) is as follows:



+-----+

The Type and Length fields are each 8 bits long. The Length field indicates the size in bytes of the remainder of the object (i.e. is the size of the address in bytes plus 2). When the address is IPv4 the length field is thus 6 and when the address is IPv6 the length field is thus 18. The length field therefore acts as both the TLV parsing parameter and the address family type indicator.

The UDP Return Object Type (URO TLV Type) has a value of 131.

The UDP Destination Port is a UDP Destination port as specified in [[RFC0768](#)].

The Address is either an IPv4 or an IPv6 address.

The URO MUST NOT appear in a response.

[4.](#) Theory of Operation

This document defines the UDP Return Object to enable the MPLS-PLDM Querier to specify the return path for the MPLS-PLDM reply using IP/UDP encapsulation.

When the MPLS-PLDM Response is requested out-of-band by setting the Control Code of the MPLS-PLDM Query to "Out-of-band Response Requested", and the URO is present, the responder SHOULD send the response back to Querier on the specified destination UDP port at the specified destination IP address contained in the URO.

If the URO is expected but is not present in a Query message and an MPLS-PLDM Response is requested out-of-band, the Query message MUST NOT be processed further, and if possible an "Error - Invalid Message" ([\[RFC6374\] Section 3.1](#)) SHOULD be sent to the Querier and the operator notified via the management system (see [Section 4.2](#) for further details).

[4.1.](#) Sending an MPLS-PM Query

When sending an MPLS-PLDM Query message, in addition to the rules and

procedures defined in [[RFC6374](#)]; the Control Code of the MPLS-PLDM Query MUST be set to "Out-of-band Response Requested", and a URO MUST be carried in the MPLS-PLDM Query message.

If the Querier uses the UDP port to de-multiplexing of the response for different measurement type, there MUST be a different UDP port for each measurement type (Delay, loss and delay-loss combined).

An 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-PLDM query messages as defined in [[RFC6374](#)] applies in this document. In addition, when an MPLS-PLDM Query request is received, with the Control Code of the MPLS-PLDM Query set to "Out-of-band Response Requested" with a URO present, then the

Responder SHOULD use that IP address and UDP port to send MPLS-PLDM response back to Querier.

If an Out-of-band response is requested and the Address object or the URO is missing, the Query SHOULD be dropped in the case of a unidirectional LSP. If both these TLVs are missing on a bidirectional LSP, the control code of Response message should set to 0x1C indicating "Error - Invalid Message" ([\[RFC6374\] Section 3.1](#)) and the 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.

[4.3.](#) Sending an MPLS-PM Response

As specified in [[RFC6374](#)] the MPLS-PLDM Response 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-PLDM Query message.

When the requested return path is an IP forwarding path and this method is in use, the destination IP address and UDP port MUST be

copied from the URO. 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 MPLS-PLDM 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 = URO.Address        |
|   Protocol = UDP                          |
|                                             |
+-----+
| UDP Header |
|   Source Port = As chosen by Responder     |
|   Destination Port = URO.UDP-Destination-Port |
|                                             |
+-----+
| Message as specified in RFC6374 |
|                                             |
+-----+
```

Figure 1: Response packet Format

If the return path is an 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 URO MAY be omitted from the MPLS-PLDM messages.

6. Security Considerations

The MPLS-PLDM 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 has made an early allocation of a new Optional TLV type from MPLS Loss/Delay Measurement TLV Object Registry contained within the g-ach-parameters registry set. IANA is requested to modify the description text as shown below.

Code	Description	Reference
131	UDP Return	[This]

[8.](#) Acknowledgements

We acknowledge the contribution of Joseph Chin and Rakesh Gandhi, both with Cisco Systems. We thank Loa Andersson, Eric Osborne, Mustapha Aissaoui, and Jeffrey Zhang for their review comments.

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

[9.](#) References

[9.1.](#) Normative References

- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, [RFC 768](#), August 1980.
- [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.

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.

[9.2.](#) Informative References

[RFC5921] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", [RFC 5921](#), July 2010.

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