

Internet Engineering Task Force (IETF)
Request for Comments: 7876
Category: Standards Track
ISSN: 2070-1721

S. Bryant
Independent
S. Sivabalan
S. Soni
Cisco Systems
July 2016

UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks

Abstract

RFC 6374 defines a protocol for Packet Loss and Delay Measurement for MPLS networks (MPLS-PLDM). This document specifies the procedures to be used when sending and processing out-of-band MPLS performance management Responses over an UDP/IP return path.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7876>.

Copyright Notice

Copyright (c) 2016 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	3
2. Requirements Language	3
3. Solution Overview	3
3.1. UDP Return Object	4
4. Theory of Operation	5
4.1. Sending an MPLS-PLDM Query	5
4.2. Receiving an MPLS-PLDM Query Request	5
4.3. Sending an MPLS-PLDM Response	6
4.4. Receiving an MPLS-PLDM Response	7
5. Congestion Considerations	7
6. Manageability Considerations	8
7. Security Considerations	8
8. IANA Considerations	8
9. References	9
9.1. Normative References	9
9.2. Informative References	9
Acknowledgements	10
Authors' Addresses	10

1. Introduction

This document describes how 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 needs to be forwarded internally 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 a particular node within a multi-CPU LSR.

The procedure described in this specification describes how the querier 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 an 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 an MPLS-PLDM Query, the responder SHOULD use the IP address and UDP port in the URO to reply back to the querier. The querier MAY include multiple UROs in an MPLS-PLDM Query indicating to the responder that an identical Response 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 that it is able to use (see Section 4.3).

The procedures defined in this document may be applied to both unidirectional and bidirectional Label Switched Paths (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

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., it is the size of the address in bytes plus 2). When the address is IPv4, the Length field is thus 6; 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 Type) has a value of 131.

The UDP-Destination-Port is a UDP destination port as specified in [RFC768].

The Address is either an IPv4 or an IPv6 address.

The URO MUST NOT appear in a Response and MUST be ignored if it is found to be present.

To prevent any ambiguity as to which address the responder needs to reply to, an MPLS-PLDM Query message containing a URO MUST NOT include an RFC 6374 Return Address TLV (TLV 1). Additionally, the method of constructing the return address from the Source Address TLV (TLV 130) described in Section 3.5.2 of [RFC6374] MUST NOT be used to construct a Response to a Query message that contains a URO.

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 UDP/IP 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-PLDM 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-multiplex the Response for a 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 the same measurement type to direct the Response to the correct management process in the LSR.

4.2. Receiving an MPLS-PLDM Query Request

The processing of MPLS-PLDM Query messages as defined in [RFC6374] applies in this document. In addition, when an MPLS-PLDM Query message 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 an MPLS-PLDM Response back to the querier.

If an out-of-band Response is requested and the URO is missing, the Query SHOULD be dropped in the case of a unidirectional LSP. If the TLV is missing on a bidirectional LSP, the Control Code of the 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 malformed request SHOULD be reported to the operator through the management system, with normal precautions taken in respect to the prevention of overload of the error-reporting system.

4.3. Sending an MPLS-PLDM Response

As specified in [RFC6374], the MPLS-PLDM Response can be sent either over 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 are copied from the URO. The source IP address and the source UDP port of the Response packet are left to the discretion of the responder, subject to the normal management and security considerations. If the querier has included URO(s) for only one IP address family and a return path of that type is not available, then the Query message MUST be discarded, and the operator SHOULD be informed of the error through the management system using the normal rate-limited approach. If the responder is configured to only respond with a single Response, and a path using the IP address family in the first URO is not available, the responder MAY search the UROs for the first URO specifying a return address family for which it does have a path and use the parameters in that URO to respond. If the responder is designed or configured not to search for a URO that it can respond to, then the operator SHOULD be informed of the error through the management system using the normal rate-limited approach.

The packet format for the MPLS-PLDM Response after the UDP header is as specified in [RFC6374]. As shown in Figure 1, the Associated 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 RFC 6374 message.

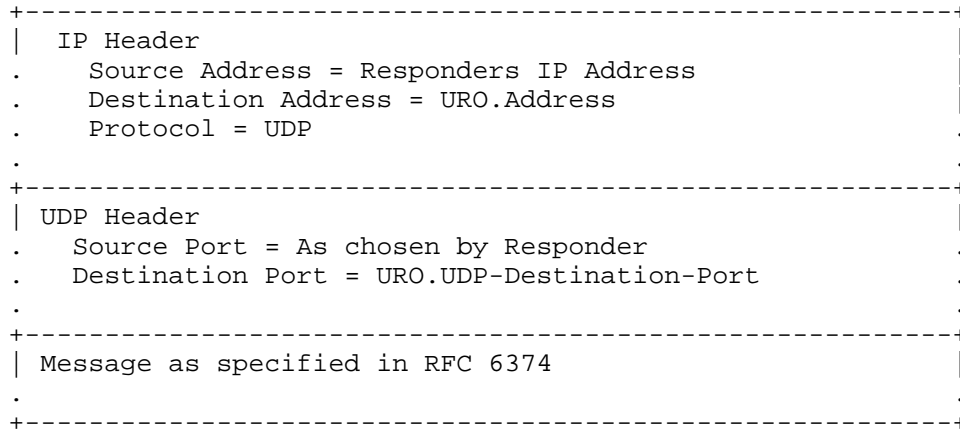


Figure 1: Response Packet Format

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-PLDM 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 reported to the operator through the management system, with normal precautions taken in respect to the prevention of overload of the error-reporting system.

5. Congestion Considerations

This protocol MUST be run in accordance with the guidance provided in [RFC5405]. As advised in Section 3.1.2 of RFC 5405, operators that wish to run this protocol at rates in excess of one packet per three seconds need to ensure that the MPLS path being monitored and any IP path that may be used to carry the Response are provisioned such that there is a negligible chance of this protocol causing congestion. Additionally, if a significant number of Response packets are lost, the querier MUST reduce the sending rate to a point where there is a negligible chance that this protocol is contributing to network congestion. The operator should also take precautions that Response

packets do not leak out of the network domain being used and cause congestion elsewhere. If a default IP address is configured by the equipment vendor, this MUST be an address known to contain the Response packet within the responder. A responder receiving a Query specifying this as a return address, and not being configured to expect such a return address, SHOULD notify the operator in a suitably rate-limited manner.

6. 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 in 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 signaling. In these circumstances, the URO MAY be omitted from the MPLS-PLDM messages.

7. 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 particular attention should be paid to the last two paragraphs. Cryptographic measures may be enhanced by the correct configuration of access-control lists and firewalls.

To prevent the use of this protocol as a reflection attack vector, the operator should ensure that the IP address in the URO addresses a system that is expecting to act as a receiver of PLDM Responses.

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

8. IANA Considerations

IANA has assigned a new optional TLV type from the "MPLS Loss/Delay Measurement TLV Object" registry contained within the "Generic Associated Channel (G-ACh) Parameters" registry set:

Code	Description	Reference
131	UDP Return	[RFC7876]

9. References

9.1. Normative References

- [RFC768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<http://www.rfc-editor.org/info/rfc768>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, DOI 10.17487/RFC3945, October 2004, <<http://www.rfc-editor.org/info/rfc3945>>.
- [RFC5405] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", BCP 145, RFC 5405, DOI 10.17487/RFC5405, November 2008, <<http://www.rfc-editor.org/info/rfc5405>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, DOI 10.17487/RFC5586, June 2009, <<http://www.rfc-editor.org/info/rfc5586>>.
- [RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, DOI 10.17487/RFC5654, September 2009, <<http://www.rfc-editor.org/info/rfc5654>>.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <<http://www.rfc-editor.org/info/rfc6374>>.

9.2. Informative References

- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, DOI 10.17487/RFC5921, July 2010, <<http://www.rfc-editor.org/info/rfc5921>>.

Acknowledgements

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

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

Authors' Addresses

Stewart Bryant
Independent

Email: stewart.bryant@gmail.com

Siva Sivabalan
Cisco Systems

Email: msiva@cisco.com

Sagar Soni
Cisco Systems

Email: sagsoni@cisco.com

