

MPLS Working Group
Internet Draft
Intended status: Standard Track
Expires: February 2011

L. Dunbar
Huawei
N. So
Verizon
P. Niger, Y. Le Goff
France Telecom
August 13, 2010

Detecting MPLS Path Impairment using MPLS-Ping

[draft-dunbar-so-mpls-detect-impair-mplsping-02.txt](#)

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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/ietf/1id-abstracts.txt>

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

This Internet-Draft will expire on February 13, 2009.

Copyright Notice

Copyright (c) 2010 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 BSD License.

Abstract

The MPLS-Ping is for detecting data path failure. This draft suggests an extension to MPLS-Ping so that transit LSR can indicate the downstream link impairment condition to the source LSR.

Conventions used in this document

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](#) 0.

Table of Contents

1.	Motivation.....	2
2.	MPLS-Ping extension for path condition impairment indication..	3
2.1.	Downstream Link Condition sub-TLV.....	4
2.2.	Link Condition Request in Echo Request.....	5
2.3.	Receiving Echo Request with a Path Condition Request flag set	6
3.	Receiving Path Condition Impairment Notification.....	6
4.	Manageability Considerations.....	7
5.	Security Considerations.....	7
6.	IANA Considerations.....	7
7.	Acknowledgments.....	7
8.	References.....	7
8.1.	Normative References.....	7
	Authors' Addresses.....	8
	Intellectual Property Statement.....	8
	Disclaimer of Validity.....	9

[1.](#) Motivation

MPLS-Ping [[MPLS-Ping](#)] can be used to detect faults along the LSP path. One of the faults detected by MPLS-Ping is downstream link

failure, i.e. link connectivity being down. In some network environment, source node also needs to know the condition of the link on which the LSPs are carried even though the link connectivity is up. That way the source LSR can perform needed functions, such as enforcing admission control, re-signal and/or re-compute the LSP

path, or generate more stringent performance monitoring at shorter interval, and etc.

A common example for the link condition change is the bandwidth fluctuation in Mobile Backhaul network, where Microwave transport is widely deployed. Most Microwave transport nodes adjust its bandwidth based on the weather. Even though there is RSVP-TE for individual links to advertise its available bandwidth in the routing domain, end-to-end bandwidth change may not be possible in some Mobile Backhaul environment, because there may be multiple routing domains from base stations to MSO. If Source Nodes, i.e. LTE's eNodeB or MSO's RNC, are aware of the bandwidth change, they can adjust services accordingly, request other base stations to accept new calls, or trigger a new Performance Monitoring scheme to track the condition more closely.

In another application, source LSRs may want to be aware of the congestion along the LSP path, so that proper actions can be taken. MPLS-ECN ([RFC 5129](#)) specifies a mechanism for transit nodes to mark EXP bits when congestion happens. However, many deployed MPLS networks already use EXP bits to mark packet priorities, making MPLS-ECN ([RFC 5129](#)) mechanism un-usable for the purpose of LSP change indication.

In a third application, source LSRs may want to be aware of significant performance degradation on the downstream links along the LSP path. The performance degradation can be increased latency and/or increased delay variation. Those performance degradations may be induced by the physical layer protection scheme, such as link switching from active side of the ring to protect side of the ring, or it may be induced by transmission media degradation.

In a forth application, source LSRs may want to be aware of transport media change on the downstream links along the LSP path. For example, the link can be a fiber protected by microwave, and the source router may be carrying an application that cannot use

This draft suggests adding a new sub-TLV to the MPLS Ping Echo Reply for the transit LSR to indicate the downstream link condition changes to source routers.

[MPLS-Ping] specified that replying router for MPLS Ping should include one Downstream Mapping for each interface, over which this FEC could be forwarded, in the Echo Reply. [MPLS-Ping-Enhanced] has

[Page 3]

August 2010

This draft suggests adding a new sub-TLV "Downstream Link Condition" to indicate the condition of the downstream link of the corresponding interface. The "Downstream Path Condition" could be bandwidth being reduced, the interface being congested, etc.

Sub-Type	Value Field
TBD	Multipath data (specified by [MPLS-Ping-Enhanced])
TBD	Label stack (specified by [MPLS-Ping-Enhanced])
TBD	FEC Stack change (specified by [MPLS-Ping-Enhanced])
TBD	Downstream Link Condition (new)

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								
Type										Length										ImpairmentType										SeverityLevel									

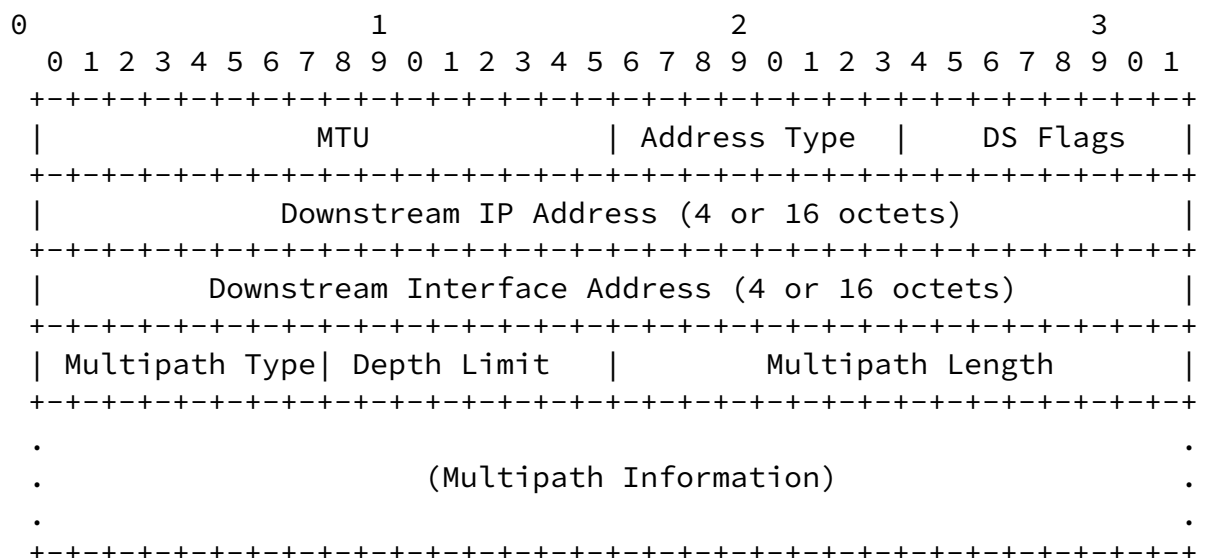
Value	Meaning
1	port towards downstream LSR is congested


- 2 Bandwidth of the link towards downstream LSR is reduced
- 3 performance of the link towards downstream LSR is reduced
- 4 transport media of the link towards downstream LSR has
 been changed

The Severity Level field is a value indicating the severity of the impairment. Network operator can set the Severity Level for anticipated conditions and configure the proper actions at the source node upon receiving the Echo Reply. For example, Severity Level 0 can represent full bandwidth, 1 represents the next bandwidth level, and 6 represents the least bandwidth. For microwave transport link within MPLS based Mobile Backhaul network, the Adaptive Modulation usually has several levels of adjusted bandwidth, which can be used as the basis for setting the Severity Level.

[2.2. Link Condition Request in Echo Request](#)

If a source LSR of a LSP cannot do anything when the LSP path is impaired, then there is no point for the transit LSR to send any link condition in the Echo Reply to the sender. Therefore, it is necessary for the sender to indicate if it desires to receive the Downstream Link Condition status in the Echo Reply. This draft suggests using one reserved bit of DS flags in the Echo Request for the source node to indicate if it desires to receive the impairment condition of the downstream link on a transit LSR.






DS Flags

[RFC4379](#) already defines I bit and N bit out of the octet DS Flags. This draft suggests adding a new bit C for source LSR to indicate if it desires to have the link impairment condition to be reported by transit LSR in the Echo Reply.

The DS flags field is a bit vector with the following format:

0 1 2 3 4 5 6 7

+---+---+---+---+---+---+---+---



I and N flags are already specified by [RFC 4379](#). The C flag is defined as follow:

Flag	Name and Meaning
C	Source LSR desires to know the downstream link condition

When this flag is set (C = 1), it indicates that the replying router SHOULD include the downstream link condition sub-TLV in the echo reply message. When this flag is not set (C=0), it indicates that the source LSR doesn't need downstream path condition information, so replying router SHOULD NOT include the downstream link condition sub-TLV in the echo reply.

[2.3](#). Receiving Echo Request with a Path Condition Request flag set

If the DS flags' C is set, the receiving LSRs HAVE TO construct the Downstream Link Condition sub-TLV and insert it into the Echo Reply.

[3.](#) Receiving Path Condition Impairment Notification

Though it is out of the scope of this draft on what Source LSR will do upon receiving the Path Condition in the Echo Reply, here are some examples of possible actions Source LSR could do:

- trigger more stringent performance monitoring in shorter interval to measure the quality of the path

- re-adjust load balancing among the multiple paths from Source LSR to the Destination LSR

- Re-signal LSP to alternative path

- activate the secondary/protection path

- adjust admission rate (in the case of LTE eNodeB)

- Notify adjacent client device via E-LMI, IEEE802.3ah, or simple flow control.

[4.](#) Manageability Considerations

This document does not add additional manageability considerations.

[5.](#) Security Considerations

This document has no additional requirement for a change to the security models of MPLS-Ping and MPLS-Ping-Enhanced.

[6.](#) IANA Considerations

A future revision of this document will present requests to IANA for codepoint allocation.

[7.](#) Acknowledgments

8. References

8.1. Normative References

- [ECN] B. Davie, et. al., "Explicit Congestion Marking in MPLS", [RFC 5129](#), Jan. 2008.
- [MPLS-PING] K. Kompella, et. al., "Detecting MPLS Data Plane Failures", [RFC 4379](#), February 2006.
- [BFD-MPLS] Aggarwal, R., "[draft-ietf-bfd-mpls-07.txt](#)", work in progress.
- [LSP-Ping-Enhanced] N. Bahadur, et. al., "Mechanism for performing LSP-Ping over MPLS tunnels", "[draft-ietf-mpls-lsp-ping-enhanced-dsmap-04](#)", work in progress.

Authors' Addresses

Linda Dunbar (Ed.)
Huawei Technologies
1700 Alma Drive, Suite 500
Plano, TX 75075, USA
Phone: (972) 543 5849
Email: ldunbar@huawei.com

Ning So (Ed.)
Enterprise Data Network and Traffic Planning
2400 N. Glenville Dr.
Richardson, TX 75082, USA

Phone: (972) 729 7905
Email: ning.so@verizonbusiness.com

Philippe Niger
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: philippe.niger@orange-ftgroup.com

Yannick Le Goff
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: yannick1.legoff@orange-ftgroup.com

Intellectual Property Statement

The IETF Trust 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 any IETF 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.

Copies of Intellectual Property 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 any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

All IETF Documents and the information contained therein 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 THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.