

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 24, 2014

Y. Yin
S. Jiang, Ed.
Z. Li
M. Chen
Huawei Technologies Co., Ltd
October 21, 2013

A Traceroute Framework in IP/MPLS Heterogeneous Networks
draft-yin-traceroute-ipmpls-00

Abstract

This document introduces a traceroute framework that can obtain information of a real traffic flow path through heterogeneous IP/MPLS network environments.

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 April 24, 2014.

Copyright Notice

Copyright (c) 2013 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.

Internet-Draft Traceroute in IP/MPLS Heterogeneous Env October 2013

Table of Contents

1.	Introduction	2
2.	Procedures of the IP/MPLS Heterogeneous Traceroute Mechanism	2
2.1.	New Components Needed	4
2.2.	Form a New Echo-Request Message	5
2.3.	Process an Echo-Request Message	5
2.3.1.	Response requested information	5
2.3.2.	Update Routing-Decide field	6
2.3.3.	Update Return-Path field	6
2.3.4.	Send a new Echo-Request message	6
2.4.	Process an Echo-Reply Message	6
2.5.	Termination of the Traceroute Request	7
3.	Security Considerations	7
4.	IANA Considerations	7
5.	Acknowledgements	7
6.	Change log [RFC Editor: Please remove]	7
7.	Informative References	8
	Authors' Addresses	8

[1.](#) Introduction

The current ISP networks may actually be constructed by IP/MPLS heterogeneous network technologies. ISP services may be delivered across these IP/MPLS heterogeneous network environments.

To locate the network fault quickly, IETF has defined the corresponding ping and trace functions for each network technology, independently. However, none of these technologies are able to gather the trace information of all these heterogeneous network environments. It is difficult to trace the complete end-to-end paths. Another issue of these ping/trace functions is path replicability. [[I-D.yin-tsvwg-ipflow-pathtrace-ps](#)] describes the issues and the requirements for a new IP/MPLS traffic flow path trace mechanism in details.

In order to meet these requirements, this document introduces a new traceroute framework that traces the real traffic flow path while the real forwarding-relevant information are carried and updated in the path. It can therefore obtain information of a real traffic flow path through IP/MPLS heterogeneous network environments.

It is worthy of noticing that this mechanism requires support of all

forwarding devices. It is suitable to be used within a single administration domain.

2. Procedures of the IP/MPLS Heterogeneous Traceroute Mechanism

Yin, et al.

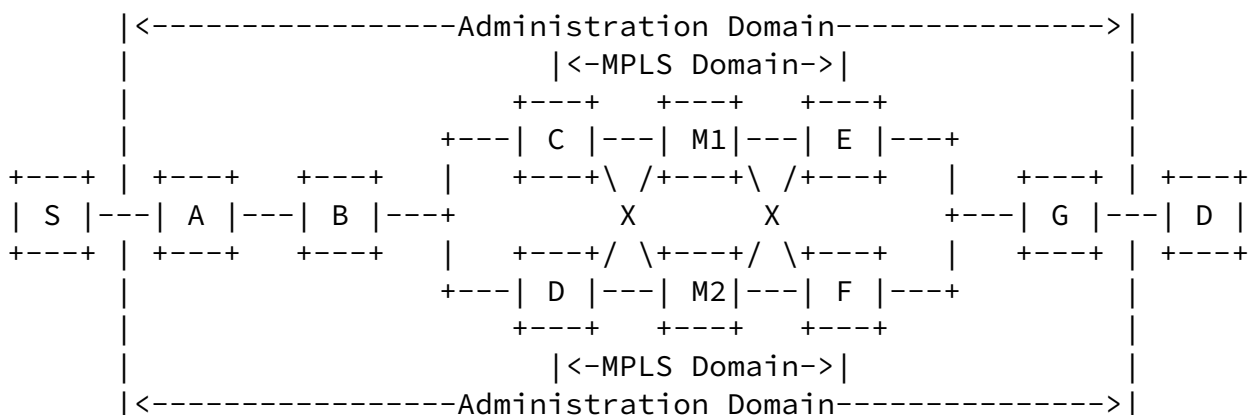
Expires April 24, 2014

[Page 2]

Internet-Draft Traceroute in IP/MPLS Heterogeneous Env October 2013


The new traceroute mechanism, introduced by this document, mainly targets the IP/MPLS heterogeneous networks.

In order to describe the applicable scenarios and procedures, in this document, below network topology illustrates a traceroute example scenario, in which the IP load balancing and traversing a MPLS network are demonstrated.



The real traffic is from the S(ource) node to the D(estination) node. The traceroute request is initialized at node A, and ideally terminated at node G. From the node B, the IP packet may be forwarded to node C or node D. Then it is MPLS domain, where the intermediate device M1 or M2 cannot be detected by current IP ping/trace mechanism. This mechanism assumes that all MPLS nodes have IP connectivities and can communicate with each other according to IP addresses.

In the new traceroute mechanism, the traceroute requesting node passes three types of information, as illustrated in the below figure, to the next hop. An example of traceroute request message:



```

+-----+
|Forwarding-Decide field contains Real IP/MPLS headers|
+-----+
|                                Options                                |
+-----+

```

The IP and UDP headers indicate the recipient - this is a traceroute request. The Forwarding-Decide field contains real IP/MPLS headers, copied from a packet in a real traffic flow, or pseudo headers forming to represent a traffic flow. It may include IP header, MPLS header and TCP/UDP header. The IP/MPLS headers in the Forwarding-

Decide field are maintained or modified in the exactly same way when a real traffic packet is processed in the forwarding procedure. Every on-path node decides the next hop according to the Forwarding-Decide field. This would guarantee the traced route follows the same path of the real traffic packet, because the forwarding decide information is exactly same. Options are mainly used to indicate the on-path nodes requested information and the return path. The on-path nodes report the requested information of themselves to the traceroute requesting node directly.

[2.1.](#) New Components Needed

There are a few new components needed in order to fulfil the newly-introduced hop-by-hop traceroute mechanism. They are listed below.

Traceroute Port A well-known UDP port that indicates the recipient devices that received UDP messages are for traceroute purpose.

An echo request message: a new UDP message that is carrying the traceroute request. It is processed, may be modified, then passed on hop by hop.

- * Transaction ID, used to distinguish multiple traceroute requests initiated by the same traceroute requesting node.
- * Forwarding-Decide field, contains all information that may affect the path choice for a specific packet, including IP header, TCP/UDP header, MPLS header, and IP tunnel header. The field contents are maintained or modified in the exactly same

way when a real IP packet is processed in the forwarding procedure.

- * Required-Information option, used to indicate the on-path nodes the information desired by the traceroute requesting node.
- * Return-Path field, used to indicate the return path for the responded Echo-Reply message. It operates like a stack: the closed relay nodes first, then second close, till the original requesting node.

An echo reply message: a new UDP message that is used to return the requested information of intermedia node to the traceroute requesting node.

- * Transaction ID, copied from the correspondent Echo-Request message. It is used to distinguish mutliples traceroute requests initiated by the same traceroute requesting node.

- * Information options, the information of the requested node. They are the response to the required information option.
- * Return-Path field, copied from the correspondent Echo-Request message. The MPLS or tunnel ingress nodes need these information to relay the Echo-Reply message to the traceroute requesting node.

The detailed format or data structure is out of scope for this informational document.

[2.2.](#) Form a New Echo-Request Message

In the new traceroute mechanism, the traceroute request device (node A) forms an Echo-Request message, in which the Forwarding-Decide field contains a real IP header and TCP/UDP header that is copied from a packet in a real traffic flow, or a pseudo IP header and TCP/UDP header, towards node D. This Echo-Request message also indicates what information of these on-path devices are desired. The traceroute request device also records itself in the Return-Path field.

The Echo-Request message is carried by an IP packet, which the destination address is filled by the loopback address and source address is the requesting node itself. IP TTL is set to value that indicate the maximum trace route hops. It then sends the new IP packet towards the next hop according to the Forwarding-Decide field.

This mechanism description in this section currently assumes that the traceroute requests start from IP.

[2.3.](#) Process an Echo-Request Message

Upon receiving the IP packet that contains an Echo-Request message, the recipient node processes it to the new traceroute function, because of the loopback address and the newly defined Traceroute port. There are four follow-up procedures on the recipient: responding its own information, updating Forwarding-Decide field, updating Return-Path field, and sending a new Echo-Request message.

[2.3.1.](#) Response requested information

The new traceroute function returns the requested information of this node, according to Required-Information option in the Echo-Request message, back to the traceroute requesting node (or the closest relay node), by sending an Echo-Reply message.

The IP address of the traceroute requesting node or the closest relay node, is obtained from the Return-Path field of the Echo-Request message. Transaction ID is copied from the Echo-Request message. The Return-Path field from the Echo-Request message must be copied into the Echo-Reply message.

[2.3.2.](#) Update Routing-Decide field

The Forwarding-Decide field is maintained/modified in the exactly same way when the real traffic packet that the Forwarding-Decide field represents is processed, including add/remove new MPLS header, update MPLS label, add/remove IP tunnel header, NAT44 or NAT66 translation or etc.

[2.3.3.](#) Update Return-Path field

The node first check whether the second close node in the Return-Path field is reachable or not. If yes, it removes the closest node in the Return-Path field; if not, do nothing. Note, if there is only one node in the Return-Path field, it must not be removed.

The node then adds the IP address of itself into the Return-Path field, and the necessary information that this node must have in order to distinguish the previous node, such as Virtual Routing and Forwarding (VRF) information, etc.

[2.3.4.](#) Send a new Echo-Request message

This node then forms a new IP packet, in which the destination address is filled by the loopback address, source address is itself. It carries the Echo-Request message in the payload. TTL in IP header is reduced by 1 every hop.

According to the MPLS header, if there is, IP header and TCP/UDP header in the Forwarding-Decide field, the device decides the next hop, then sends the new IP packet towards the next hop.

[2.4.](#) Process an Echo-Reply Message

Upon receiving an Echo-Reply message that the destination IP address is itself, the recipient must decide whether the Echo-Reply is terminate here or needs to be relayed out, by checking whether it is the last node in the Return-Path field.

If it is not the last node, it must be the first node. It then obtains the necessary information for it to distinguish the next node. It removes itself out of the stack of the Return-Path field. It then relays the Echo-Reply message to the next relay node or the original requesting node that it learns from the Return-Path field.

[2.5.](#) Termination of the Traceroute Request

Ideally, the traceroute request should be terminated at the edge of the administration domain if the traffic destination was out of domain, Node G in the example scenario. The forwarding devices on the edge of the administration domain should be configured not to send traceroute messages out on the interfaces that connected to subscribers or devices out of the administration domain.

However, the leak of traceroute request message should not bring much impact. The subscriber node does not support this traceroute function would drop it silently as unknown message. The ingress devices of another administration domain or another ISP/transit network would filter this message.

[3.](#) Security Considerations

Without an authentication mechanism, this mechanism would be risk to expose network device information. It should only be used either when combines with an authentication mechanism or in a closed single administration domain, in which the end user requests or request from outside of this administration domain should be filtered at the edge of the administration domain.

The leak of traceroute request message does not create much security risk. The information carried with in the message does not contain much sensitive information of the network. The only potential risk information is exposing of the IP addresses of intermediate forwarding devices in the return path option.

[4.](#) IANA Considerations

This memo includes no request to IANA.

[5.](#) Acknowledgements

This document was produced using the xml2rfc tool [[RFC2629](#)].

[6.](#) Change log [RFC Editor: Please remove]

[draft-yin-hopbyhop-traceroute-00](#): original version. 2013-10-21.

[7.](#) Informative References

[I-D.yin-tsvwg-ipflow-pathtrace-ps]

Yin, Y., Jiang, S., and G. Yan, "IP Flow Path Trace Requirements", [draft-yin-tsvwg-ipflow-pathtrace-ps-00](#) (work in progress), July 2013.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#), June 1999.

Authors' Addresses

Yuanbin Yin
Huawei Technologies Co., Ltd
Q15, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: yinyuanbin@huawei.com

Sheng Jiang (editor)
Huawei Technologies Co., Ltd
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: jiangsheng@huawei.com

Zhenbin Li
Huawei Technologies Co., Ltd
Q15, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: lizhenbin@huawei.com

Mach(Guoyi) Chen
Huawei Technologies Co., Ltd
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: mach.chen@huawei.com