Workgroup: MPLS WG
Internet-Draft:
draft-ietf-mpls-lspping-norao-08
Updates: 8029 (if approved)
Published: 1 March 2024
Intended Status: Standards Track
Expires: 2 September 2024
Authors: K. Kompella (Juniper Networks)
R. Bonica (Juniper Networks)
G. Mirsky, Ed. (Ericsson)

Deprecating the Use of Router Alert in LSP Ping

Abstract

The MPLS echo request and MPLS echo response messages, defined in RFC 8029 "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures" (usually referred to as LSP ping messages), are encapsulated in IP whose headers include a Router Alert Option (RAO). In actual deployments, the RAO was neither required nor used. Furthermore, RFC 6398 identifies security vulnerabilities associated with the RAO in non-controlled environments, e.g., the case of using the MPLS echo request/reply as inter-area Operations, Administration, and Maintenance (OAM), and recommends against its use outside of controlled environments.

Therefore, this document retires the RAO for MPLS OAM and updates RFC 8029 to remove the RAO from LSP ping message encapsulations. Furthermore, this document explains why RFC 7506 has been reclassified as Historic.

Also, the use of an IPv6 loopback address (::1/128) as the IPv6 destination address for an MPLS echo request message is RECOMMENDED.

Table of Contents

1. Note for the RFC Editor

Per IESG decision, this document MUST be processed only after the status of RFC 7506 is changed to Historical. This note must be removed before the publication.

2. Introduction

RFC 8029 - "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures" (usually referred to as LSP Ping) [RFC8029] detects data-plane failures in MPLS Label Switched Paths (LSPs). It can operate in "ping mode" or "traceroute mode". When operating in ping mode, it checks LSP connectivity. When operating in traceroute mode, it can trace an LSP and localize failures to a particular node along an LSP.

The reader is assumed be familiar with [RFC8029] and its terminology.

LSP ping defines a probe message called the "MPLS echo request". It also defines a response message called the "MPLS echo reply". Both messages are encapsulated in UDP and IP. The MPLS echo request message is further encapsulated in an MPLS label stack, except when all of the Forwarding Equivalency Classes in the stack correspond to Implicit Null labels.

When operating in ping mode, LSP ping sends a single MPLS echo request message, with the MPLS TTL set to 255. This message is intended to reach the egress Label Switching Router (LSR). When operating in traceroute mode, MPLS ping sends multiple MPLS echo request messages as defined in Section 4.3 of [RFC8029]. It manipulates the MPLS TTL so that the first message expires on the first LSR along the path and subsequent messages expire on subsequent LSRs.

According to [RFC8029], the IP header that encapsulates an MPLS echo request message must include a Router Alert Option (RAO). Furthermore, [RFC8029] also says that the IP header that encapsulates an MPLS echo reply message must include an RAO if the value of the Reply Mode in the corresponding MPLS echo request message is "Reply via an IPv4/IPv6 UDP packet with Router Alert". This document explains why RAO was not needed in both cases. Furthermore, [RFC6398] identifies security vulnerabilities associated with the RAO in non-controlled environments, e.g., the case of using the MPLS echo request/reply as inter-domain OAM over the public Internet, and recommends against its use outside of controlled environments, e.g., outside a single administrative domain.

Therefore, this document updates RFC 8029 [RFC8029] to retire the RAO from both LSP ping message encapsulations and explains why RFC 7506 [RFC7506] has been reclassified as Historic.

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Router Alert for LSP Ping (RFC 8029)

3.1. MPLS Echo Request

While the MPLS echo request message must traverse every node in the LSP under test, it must not traverse any other node. Specifically, the message must not be forwarded beyond the egress Label Switching Router (LSR). To achieve this, a set of the mechanisms that are used

concurrently to prevent leaking MPLS echo request messages has been defined in [RFC8029]:

1. When the MPLS echo request message is encapsulated in IPv4, the IPv4 destination address must be chosen from the subnet 127/8. When the MPLS echo request message is encapsulated in IPv6, the IPv6 destination address must be chosen from the subnet 0:0:0:0:0:FFFF:7F00:0/104.

2. When the MPLS echo request message is encapsulated in IPv4, the IPv4 TTL must be equal to 1. When the MPLS echo request message is encapsulated in IPv6, the IPv6 Hop Limit must be equal to 1. For further information on the encoding of the TTL/Hop Limit in an MPLS echo request message, see Section 4.3 of [RFC8029].

3. When the MPLS echo request message is encapsulated in IPv4, the IPv4 header must include an RAO with the option value set to "Router shall examine packet" [RFC2113]. When the MPLS echo request message is encapsulated in IPv6, the IPv6 header chain must include a Hop-by-hop extension header and the Hop-by-hop extension header must include an RAO with the option value set to MPLS OAM [RFC7506].

Currently, all of these are required. However, any one is sufficient to prevent forwarding the packet beyond the egress LSR.

Therefore, this document updates RFC 8029 [RFC8029] in that Requirement 3 is removed.

No implementation that relies on the RAO to prevent packets from being forwarded beyond the egress LSR have been reported to the MPLS working group.

3.2. MPLS Echo Reply

An LSP ping replies to the MPLS echo request message with an MPLS echo reply message. Four reply modes are defined in [RFC8029]:

1. Do not reply

2. Reply via an IPv4/IPv6 UDP packet

3. Reply via an IPv4/IPv6 UDP packet with Router Alert

4. Reply via application-level control channel

The rationale for mode 3 is questionable, if not wholly misguided. According to RFC 8029 [RFC8029], "If the normal IP return path is deemed unreliable, one may use 3 (Reply via an IPv4/IPv6 UDP packet with Router Alert)."

However, it is not clear that the use of the RAO increases the reliability of the return path. In fact, one can argue it decreases the reliability in many instances, due to the additional burden of processing the RAO. This document updates RFC 8029 [RFC8029] in that mode 3 is removed.

No implementations of mode 3 have been reported to the MPLS working group.

4. Reclassification of RFC 7506 as Historic

RFC 7506 [RFC7506] defines the IPv6 Router Alert Option for MPLS Operations, Administration, and Management. This document explains why RFC 7506 [RFC7506] has been reclassified as Historic.

5. Update to RFC 8029

[RFC8029] requires that the IPv6 Destination Address used in IP/UDP encapsulation of an MPLS echo request packet is selected from the IPv4 loopback address range mapped to IPv6. Such packets do not have the same behavior as prescribed in [RFC1122] for an IPv4 loopback addressed packet.

[RFC4291] defines ::1/128 as the single IPv6 loopback address. Considering that, this specification updates Section 2.1 of [RFC8029] regarding the selection of an IPv6 destination address for an MPLS echo request message as follows:

OLD

The 127/8 range for IPv4 and that same range embedded in an IPv4-mapped IPv6 address for IPv6 was chosen for a number of reasons.

RFC 1122 allocates the 127/8 as the "Internal host loopback address" and states: "Addresses of this form MUST NOT appear outside a host." Thus, the default behavior of hosts is to discard such packets. This helps to ensure that if a diagnostic packet is misdirected to a host, it will be silently discarded.

RFC 1812 [RFC1812] states:

> *A router SHOULD NOT forward, except over a loopback interface, any packet that has a destination address on network 127. A router MAY have a switch that allows the network manager to disable these checks. If such a switch is provided, it MUST default to performing the checks.

This helps to ensure that diagnostic packets are never IP forwarded.

The 127/8 address range provides 16M addresses allowing wide flexibility in varying addresses to exercise ECMP paths. Finally, as an implementation optimization, the 127/8 range provides an easy means of identifying possible LSP packets.

NEW

The 127/8 range for IPv4 was chosen for a number of reasons.

RFC 1122 [RFC1122] allocates the 127/8 as the "Internal host loopback address" and states: "Addresses of this form MUST NOT appear outside a host." Thus, the default behavior of hosts is to discard such packets. This helps to ensure that if a diagnostic packet is misdirected to a host, it will be silently discarded.

RFC 1812 [RFC1812] states:

> *A router SHOULD NOT forward, except over a loopback interface, any packet that has a destination address on network 127. A router MAY have a switch that allows the network manager to disable these checks. If such a switch is provided, it MUST default to performing the checks.

This helps to ensure that diagnostic packets are never IP forwarded.

The 127/8 address range provides 16M addresses allowing wide flexibility in varying addresses to exercise ECMP paths. Finally, as an implementation optimization, the 127/8 range provides an easy means of identifying possible LSP packets.

The IPv6 destination address for an MPLS echo request message is selected as follows:

*The IPv6 loopback address ::1/128 SHOULD be used.

*The sender of an MPLS echo request MAY select the IPv6 destination address from the 0:0:0:0:0:FFFF:7F00/104 range.

*To exercise all paths in an ECMP environment, the source of entropy other than the IP destination address SHOULD be used. For example, MPLS Entropy Label [RFC6790] or IPv6 Flow Label [RFC6438] can be used as the source of entropy.

END

Additionally, this specification updates Section 2.2 of [RFC8029] to replace the whole of the section with the following text:

> LSP Ping implementations SHOULD ignore RAO options when they arrive on incoming MPLS echo request and MPLS echo reply messages.

Resulting from the removal of the Reply mode 3 "Reply via an IPv4/IPv6 UDP packet with Router Alert" (see Section 3.2), this specification updates Section 4.5 of [RFC8029] by removing the following text:

> If the Reply Mode in the echo request is "Reply via an IPv4 UDP packet with Router Alert", then the IP header MUST contain the Router Alert IP Option of value 0x0 [RFC2113] for IPv4 or 69 [RFC7506] for IPv6. If the reply is sent over an LSP, the topmost label MUST in this case be the Router Alert label (1) (see [RFC3032]).

Furthermore, this specification updates Section 4.3 of [RFC8029] as follows:

OLD:

The Router Alert IP Option of value 0x0 [RFC2113] for IPv4 or value 69 [RFC7506] for IPv6 MUST be set in the IP header.

NEW:

The Router Alert IP Option of value 0x0 [RFC2113] for IPv4 or value 69 [RFC7506] for IPv6 MUST NOT be set in the IP header.

END

6. Backwards Compatibility

LSP Ping implementations that conform to this specification SHOULD ignore RAO options when they arrive on incoming MPLS echo request and MPLS echo reply messages. However, this will not harm backwards compatibility because other mechanisms will also be in use by all legacy implementations in the messages they send and receive.

Section 7 of this document deprecates the IPv6 RAO value for MPLS OAM (69) in [IANA-IPV6-RAO] and the Reply Mode 3 ("Reply via an IPv4/IPv6 UDP packet with Router Alert") in [IANA-LSP-PING].

[RFC8126] offers a formal description of the word "Deprecated". In this context, "Deprecated" means that the deprecated values SHOULD NOT be used in new implementations, and that deployed implementations that already use these values continue to work seamlessly.

7. IANA Considerations

IANA is requested to mark the IPv6 RAO value of MPLS OAM (69) in [IANA-IPV6-RAO] as "Deprecated".

IANA is also requested to mark Reply Mode 3 ("Reply via an IPv4/IPv6 UDP packet with Router Alert") in "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters"[IANA-LSP-PING] as "Deprecated".

8. Security Considerations

The recommendations this document makes do not compromise security. In case of using IPv6 loopback address ::1/128 strengthens security for LSP Ping by using the standardized loopback address with well-defined behavior.

9. Acknowledgments

The authors express their appreciation to Adrian Farrel and Gyan Mishra for their suggestions that improved the readability of the document.

10. Normative References

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/info/rfc1122>.

[RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <https://www.rfc-editor.org/info/rfc1812>.

[RFC2113] Katz, D., "IP Router Alert Option", RFC 2113, DOI 10.17487/RFC2113, February 1997, <https://www.rfc-editor.org/info/rfc2113>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <https://www.rfc-editor.org/info/rfc6398>.

[RFC7506] Raza, K., Akiya, N., and C. Pignataro, "IPv6 Router Alert Option for MPLS Operations, Administration, and Maintenance (OAM)", RFC 7506, DOI 10.17487/RFC7506, April 2015, <https://www.rfc-editor.org/info/rfc7506>.

[RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <https://www.rfc-editor.org/info/rfc8029>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

11. Informational References

[IANA-IPV6-RAO] IANA, "IPv6 Router Alert Option Values", n.d., <https://www.iana.org/assignments/ipv6-routeralert-values>.

[IANA-LSP-PING] IANA, "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters", n.d., <https://www.iana.org/assignments/mpls-lsp-ping-parameters/mpls-lsp-ping-parameters.xml>.

[RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <https://www.rfc-editor.org/info/rfc6438>.

[RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <https://www.rfc-editor.org/info/rfc6790>.

Authors' Addresses

Kireeti Kompella
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
United States

Email: kireeti.ietf@gmail.com

Ron Bonica
Juniper Networks
1133 Innovation Way

Sunnyvale, CA 94089
United States

Email: rbonica@juniper.net

Greg Mirsky (editor)
Ericsson

Email: gregimirsky@gmail.com