

Internet Engineering Task Force (IETF)
Request for Comments: 7746
Category: Standards Track
ISSN: 2070-1721

R. Bonica
Juniper Networks
I. Minei
Google, Inc.
M. Conn
D. Pacella
L. Tomotaki
Verizon
January 2016

Label Switched Path (LSP) Self-Ping

Abstract

When certain RSVP-TE optimizations are implemented, ingress Label Switching Router (LSRs) can receive RSVP RESV messages before forwarding state has been installed on all downstream nodes. According to the RSVP-TE specification, the ingress LSR can forward traffic through a Label Switched Path (LSP) as soon as it receives a RESV message. However, if the ingress LSR forwards traffic through the LSP before forwarding state has been installed on all downstream nodes, traffic can be lost.

This document describes LSP Self-ping. When an ingress LSR receives an RESV message, it can invoke LSP Self-ping procedures to ensure that forwarding state has been installed on all downstream nodes.

LSP Self-ping is a new protocol. It is not an extension of LSP Ping. Although LSP Ping and LSP Self-ping are named similarly, each is designed for a unique purpose. Each protocol listens on its own UDP port and executes its own procedures.

LSP Self-ping is an extremely lightweight mechanism. It does not consume control-plane resources on transit or egress LSRs.

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 5741.

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/rfc7746>.

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
1.1. Requirements Language	4
2. Applicability	4
3. The LSP Self-ping Message	6
4. LSP Self-Ping Procedures	7
5. Bidirectional LSP Procedures	8
6. IANA Considerations	9
7. Security Considerations	9
8. References	9
8.1. Normative References	9
8.2. Informative References	10
Appendix A. Rejected Approaches	11
Acknowledgements	11
Contributors	12
Authors' Addresses	12

1. Introduction

Ingress Label Switching Routers (LSRs) use RSVP-TE [RFC3209] to establish MPLS Label Switched Paths (LSPs). The following paragraphs describe RSVP-TE procedures.

The ingress LSR calculates a path between itself and an egress LSR. The calculated path can be either strictly or loosely routed. Having calculated a path, the ingress LSR constructs an RSVP PATH message. The PATH message includes an Explicit Route Object (ERO) that represents the path between the ingress and egress LSRs.

The ingress LSR forwards the PATH message towards the egress LSR, following the path defined by the ERO. Each transit LSR that receives the PATH message executes admission control procedures. If the transit LSR admits the LSP, it sends the PATH message downstream, to the next node in the ERO.

When the egress LSR receives the PATH message, it binds a label to the LSP. The label can be implicit null, explicit null, or non-null. The egress LSR then installs forwarding state (if necessary) and constructs an RSVP RESV message. The RESV message contains a Label Object that includes the label that has been bound to the LSP.

The egress LSR sends the RESV message upstream towards the ingress LSR. The RESV message visits the same transit LSRs that the PATH message visited, in reverse order. Each transit LSR binds a label to the LSP, updates its forwarding state, and updates the RESV message. As a result, the Label Object in the RESV message contains the label that has been bound to the LSP most recently. Finally, the transit LSR sends the RESV message upstream, along the reverse path of the LSP.

When the ingress LSR receives the RESV message, it installs forwarding state. Once the ingress LSR installs forwarding state, it can forward traffic through the LSP.

Referring to any LSR, RFC 3209 says, "The node SHOULD be prepared to forward packets carrying the assigned label prior to sending the Resv message." However, RFC 3209 does not strictly require this behavior.

Some implementations optimize the above-described procedure by allowing LSRs to send RESV messages before installing forwarding state [RFC6383]. This optimization is desirable, because it allows LSRs to install forwarding state in parallel, thus accelerating the process of LSP signaling and setup. However, this optimization creates a race condition. When the ingress LSR receives a RESV message, some downstream LSRs may have not yet installed forwarding

state. If the ingress LSR forwards traffic through the LSP before forwarding state has been installed on all downstream nodes, traffic can be lost.

This document describes LSP Self-ping. When an ingress LSR receives an RESV message, it can invoke LSP Self-ping procedures to verify that forwarding state has been installed on all downstream nodes. By verifying the installation of downstream forwarding state, the ingress LSR eliminates this particular cause of traffic loss.

LSP Self-ping is a new protocol. It is not an extension of LSP Ping [RFC4379]. Although LSP Ping and LSP Self-ping are named similarly, each is designed for a unique purpose. Each protocol listens on its own UDP port and executes its own procedures.

LSP Self-ping is an extremely lightweight mechanism. It does not consume control-plane resources on transit or egress LSRs.

1.1. 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].

2. Applicability

LSP Self-ping is applicable in the following scenario:

- o The ingress LSR signals a point-to-point LSP.
- o The ingress LSR receives a RESV message.
- o The RESV message indicates that all downstream nodes have begun the process of forwarding state installation.
- o The RESV message does not guarantee that all downstream nodes have completed the process of forwarding state installation.
- o The ingress LSR needs to confirm that all downstream nodes have completed the process for forwarding state installation.
- o The ingress LSR does not need to confirm the correctness of downstream forwarding state, because there is a very high likelihood that downstream forwarding state is correct.
- o Control-plane resources on the egress LSR may be scarce.

- o The need to conserve control-plane resources on the egress LSR outweighs the need to determine whether downstream forwarding state is correct.

Unlike LSP Ping and S-BFD [S-BFD], LSP Self-ping is not a general-purpose MPLS OAM mechanism. It cannot reliably determine whether downstream forwarding state is correct. For example, if a downstream LSR installs a forwarding state that causes an LSP to terminate at the wrong node, LSP Self-ping will not detect an error. In another example, if a downstream LSR erroneously forwards a packet without an MPLS label, LSP Self-ping will not detect an error.

Furthermore, LSP Self-ping fails when either of the following conditions are true:

- o The LSP under test is signaled by the Label Distribution Protocol (LDP) Independent Mode [RFC5036].
- o Reverse Path Forwarding (RPF) [RFC3704] filters are enabled on links that connect the ingress LSR to the egress LSR.

While LSP Ping and S-BFD are general-purpose OAM mechanisms, they are not applicable in the above-described scenario because:

- o LSP Ping consumes control-plane resources on the egress LSR.
- o An S-BFD implementation either consumes control-plane resources on the egress LSR or requires special support for S-BFD on the forwarding plane.

By contrast, LSP Self-ping requires nothing from the egress LSR beyond the ability to forward an IP datagram.

LSP Self-ping's purpose is to determine whether forwarding state has been installed on all downstream LSRs. Its primary constraint is to minimize its impact on egress LSR performance. This functionality is valuable during network convergence events that impact a large number of LSPs.

Therefore, LSP Self-ping is applicable in the scenario described above, where the LSP is signaled by RSVP, RPF is not enabled, and the need to conserve control-plane resources on the egress LSR outweighs the need to determine whether downstream forwarding state is correct.


3. The LSP Self-ping Message

The LSP Self-ping Message is a User Datagram Protocol (UDP) [RFC768] packet that encapsulates a session ID. If the RSVP messages used to establish the LSP under test were delivered over IPv4 [RFC791], the UDP datagram MUST be encapsulated in an IPv4 header. If the RSVP messages used to establish the LSP were delivered over IPv6 [RFC2460], the UDP datagram MUST be encapsulated in an IPv6 header.

In either case:

- o The IP Source Address MAY be configurable. By default, it MUST be the address of the egress LSR.
- o The IP Destination Address MUST be the address of the ingress LSR.
- o The IP Time to Live (TTL) / Hop Count MAY be configurable. By default, it MUST be 255.
- o The IP DSCP (Differentiated Services Code Point) MAY be configurable. By default, it MUST be CS6 (110000) [RFC4594].
- o The UDP Source Port MUST be selected from the dynamic range (49152-65535) [RFC6335].
- o The UDP Destination Port MUST be lsp-self-ping (8503) [IANA.PORTS]

UDP packet contents have the following format:



LSP Self-Ping Message

The Session-ID is a 64-bit field that associates an LSP Self-ping message with an LSP Self-ping session.

4. LSP Self-Ping Procedures

In order to verify that an LSP is ready to carry traffic, the ingress LSR creates a short-lived LSP Self-ping session. All session state is maintained locally on the ingress LSR. Session state includes the following information:

- o Session-ID: A 64-bit number that identifies the LSP Self-ping session.
- o Retry Counter: The maximum number of times that the ingress LSR probes the LSP before terminating the LSP Self-ping session. The initial value of this variable is determined by configuration.
- o Retry Timer: The number of milliseconds that the LSR waits after probing the LSP. The initial value of this variable is determined by configuration.
- o Status: A boolean variable indicating the completion status of the LSP Self-ping session. The initial value of this variable is FALSE.

Implementations MAY represent the above-mentioned information in any format that is convenient to them.

The ingress LSR executes the following procedure until Status equals TRUE or Retry Counter equals zero:

- o Format a LSP Self-ping message.
- o Set the Session-ID in the LSP Self-ping message to the Session-ID mentioned above.
- o Send the LSP Self-ping message through the LSP under test.
- o Set a timer to expire in Retry Timer milliseconds.
- o Wait until either an LSP Self-ping message associated with the session returns or the timer expires. If an LSP Self-ping message associated with the session returns, set Status to TRUE. Otherwise, decrement the Retry Counter. Optionally, increase the value of Retry Timer according to an appropriate back-off algorithm.

In the process described above, the ingress LSR addresses an LSP Self-ping message to itself and forwards that message through the LSP under test. If forwarding state has been installed on all downstream LSRs, the egress LSR receives the LSP Self-ping message and

determines that it is addressed to the ingress LSR. So, the egress LSR forwards the LSP Self-ping message back to the ingress LSR, exactly as it would forward any other IP packet.

The LSP Self-ping message can arrive at the egress LSR with or without an MPLS header, depending on whether the LSP under test executes penultimate hop-popping procedures. If the LSP Self-ping message arrives at the egress LSR with an MPLS header, the egress LSR removes that header.

If the egress LSR's most preferred route to the ingress LSR is through an LSP, the egress LSR forwards the LSP Self-ping message through that LSP. However, if the egress LSR's most preferred route to the ingress LSR is not through an LSP, the egress LSR forwards the LSP Self-ping message without MPLS encapsulation.

When an LSP Self-ping session terminates, it returns its completion status to the invoking protocol. For example, if RSVP-TE invokes LSP Self-ping as part of the LSP setup procedure, LSP Self-ping returns its completion status to RSVP-TE.

5. Bidirectional LSP Procedures

A bidirectional LSP has an active side and a passive side. The active side calculates the ERO and signals the LSP in the forward direction. The passive side reverses the ERO and signals the LSP in the reverse direction.

When LSP Self-ping is applied to a bidirectional LSP:

- o The active side calculates the ERO, signals the LSP, and runs LSP Self-ping.
- o The Passive side reverses the ERO, signals the LSP, and runs another instance of LSP Self-ping.
- o Neither side forwards traffic through the LSP until local LSP Self-ping returns TRUE.

The two LSP Self-ping sessions mentioned above are independent of one another. They are not required to have the same Session-ID. Each endpoint can forward traffic through the LSP as soon as its local LSP Self-ping returns TRUE. Endpoints are not required to wait until both LSP Self-ping sessions have returned TRUE.

6. IANA Considerations

IANA has assigned UDP Port Number 8503 [IANA.PORTS] for use by MPLS LSP Self-Ping.

7. Security Considerations

LSP Self-ping messages are easily forged. Therefore, an attacker can send the ingress LSR a forged LSP Self-ping message, causing the ingress LSR to terminate the LSP Self-ping session prematurely. In order to mitigate these threats, operators SHOULD filter LSP Self-ping packets at the edges of the MPLS signaling domain. Furthermore, implementations SHOULD NOT assign Session-IDs in a predictable manner. In order to avoid predictability, implementations can leverage a Cryptographically Secure Pseudorandom Number Generator (CSPRNG) [NIST-CSPRNG].

8. References

8.1. 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>>.
- [RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<http://www.rfc-editor.org/info/rfc791>>.
- [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>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<http://www.rfc-editor.org/info/rfc3704>>.

- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, DOI 10.17487/RFC4379, February 2006, <<http://www.rfc-editor.org/info/rfc4379>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<http://www.rfc-editor.org/info/rfc5036>>.
- [RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", BCP 165, RFC 6335, DOI 10.17487/RFC6335, August 2011, <<http://www.rfc-editor.org/info/rfc6335>>.

8.2. Informative References

- [IANA.PORTS] IANA, "Service Name and Transport Protocol Port Number Registry", <<http://www.iana.org/assignments/service-names-port-numbers>>.
- [NIST-CSPRNG] NIST, "Recommendation for Random Number Generation Using Deterministic Random Bit Generators", NIST Special Publication 800-90A, January 2012.
- [RFC4594] Babiarez, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, DOI 10.17487/RFC4594, August 2006, <<http://www.rfc-editor.org/info/rfc4594>>.
- [RFC6383] Shiomoto, K. and A. Farrel, "Advice on When It Is Safe to Start Sending Data on Label Switched Paths Established Using RSVP-TE", RFC 6383, DOI 10.17487/RFC6383, September 2011, <<http://www.rfc-editor.org/info/rfc6383>>.
- [S-BFD] Akiya, N., Pignataro, C., Ward, D., Bhatia, M., and J. Networks, "Seamless Bidirectional Forwarding Detection (S-BFD)", Work in Progress, draft-ietf-bfd-seamless-base-05, June 2015.

Appendix A. Rejected Approaches

In a rejected approach, the ingress LSR uses LSP Ping to verify LSP readiness. This approach was rejected for the following reasons.

While an ingress LSR can control its control-plane overhead due to LSP Ping, an egress LSR has no such control. This is because each ingress LSR can, on its own, control the rate of the LSP Ping originated by the LSR, while an egress LSR must respond to all the LSP Pings originated by various ingresses. Furthermore, when an MPLS Echo Request reaches an egress LSR, it is sent to the control plane of the egress LSR; this makes egress LSR processing overhead of LSP Ping well above the overhead of its data plane (MPLS/IP forwarding). These factors make LSP Ping problematic as a tool for detecting LSP readiness to carry traffic when dealing with a large number of LSPs.

By contrast, LSP Self-ping does not consume any control-plane resources at the egress LSR, and it relies solely on the data plane of the egress LSR, making it more suitable as a tool for checking LSP readiness when dealing with a large number of LSPs.

In another rejected approach, the ingress LSR does not verify LSP readiness. Instead, it sets a timer when it receives an RSVP RESV message and does not forward traffic through the LSP until the timer expires. This approach was rejected because it is impossible to determine the optimal setting for this timer. If the timer value is set too low, it does not prevent black-holing. If the timer value is set too high, it slows down the process of LSP signaling and setup.

Moreover, the above-mentioned timer is configured on a per-router basis. However, its optimum value is determined by a network-wide behavior. Therefore, changes in the network could require changes to the value of the timer, making the optimal setting of this timer a moving target.

Acknowledgements

Thanks to Yakov Rekhter, Ravi Singh, Eric Rosen, Eric Osborne, Greg Mirsky, and Nobo Akiya for their contributions to this document.

Contributors

The following individuals contributed significantly to this document:

Mark Wygant
Verizon
mark.wygant@verizon.com

Ravi Torvi
Juniper Networks
rtorvi@juniper.net

Authors' Addresses

Ron Bonica
Juniper Networks

Email: rbonica@juniper.net

Ina Minei
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
United States

Email: inaminei@google.com

Michael Conn
Verizon

Email: meconn26@gmail.com

Dante Pacella
Verizon

Email: dante.j.pacella@verizon.com

Luis Tomotaki
Verizon

Email: luis.tomotaki@verizon.com

