Network Working Group
INTERNET-DRAFT
Intended Status: Proposed Standard
Expires: July 10, 2022

J. Gao
J. Dai

Fiberhome Telecom LTD

January 10, 2022

State-updating mechanism in RSVP-TE for MPLS network
draft-gao-mpls-teas-rsvpte-state-update-04

Abstract

RSVP-TE has the following advantages: source routing capability, and the ability to reserve resources hop by hop along the LSP path. The two advantages are used by Deterministic Networking (DetNet) to provide DetNet Quality of Service (QoS) in a fully distributed control plane utilizing dynamic signaling protocols or in a Combined Control Plane (partly centralized, partly distributed).

RSVP takes a "soft state" approach to manage the reservation state in routers and hosts. The use of 'Refresh messages' to cover many possible failures has resulted in a number of operational problems. One problem relates to scaling, another relates to the reliability and latency of RSVP Signaling.

This document describes a number of mechanisms that can be used to reduce processing overhead requirements of refresh messages. These extension present no backwards compatibility issues.

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 https://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."

Table of Contents

1. Introduction

Standard RSVP [RFC2205] maintains state via the generation of RSVP refresh messages. Refresh messages are used to both synchronize state between RSVP neighbors and to recover from lost RSVP messages. The use of Refresh messages to cover many possible failures has resulted in a number of operational problems. One problem relates to scaling, another relates to the reliability and latency of RSVP Signaling.

The scaling problems are linked to the resource requirements (in terms of processing and memory) of running RSVP. The resource requirements increase proportionally with the number of sessions. Each session requires the generation, transmission, reception and processing of RSVP Path and Resv messages per refresh period. Supporting a large number of sessions, and the corresponding volume of refresh messages, presents a scaling problem.

The reliability and latency problem occurs when a non-refresh RSVP message is lost in transmission. Standard RSVP [RFC2205] recovers from a lost message via RSVP refresh messages. In the face of transmission loss of RSVP messages, the end-to-end latency of RSVP signaling is tied to the refresh interval of the node(s) experiencing the loss. When end-to-end signaling is limited by the refresh interval, the delay incurred in the establishment or the change of a reservation may be beyond the range of what is acceptable for some applications.

This document proposes to disable RSVP refresh messages to solve soft-state scaling problems. The reliable message delivery mechanism specified in [RFC2961] states that "Nodes receiving a non-out of order message containing a MESSAGE_ID object with the ACK_Desired flag set, SHOULD respond with a MESSAGE_ID_ACK object.". When RSVP refresh messages are disabled, the time to deallocate resources after a tear message is lost is an issue. To solve this problem, MUST make use of the Hello session based on the Node-ID ([RFC3209][RFC4558]) for detection of RSVP-TE signaling adjacency failures. MUST implement coupling the state of individual LSPs with the state of the corresponding RSVP-TE signaling adjacency. When an RSVP-TE speaker detects RSVP-TE signaling adjacency failure, the speaker MUST act as if all the Path and Resv states learned via the failed signaling adjacency have timed out. To avoid compatibility problems, a flag bit in the RSVP message header is extended to disable RSVP refresh messages.

2. Terminology

2.1. Terms Used in This Document

Refresh messages: represent previously advertised state and contain exactly the same objects and same information as a previously transmitted message, and are sent over the same path. Only Path and Resv messages can be refresh messages. Refresh messages are identical to the corresponding previously transmitted message, with some possible exceptions.

Trigger messages: Trigger messages are those RSVP messages that advertise state or any other information not previously transmitted. Trigger messages include messages advertising new state, a route change that alters a reservation path, or a modification to an existing RSVP session or reservation.

2.2. Abbreviations

The following abbreviations are used in this document:
RSVP: Resource ReserVation Protocol.
RSVP-TE: Resource ReserVation Protocol - Traffic Engineering.

3 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.

4 State-updating mechanism in RSVP for MPLS network

To indicate support for the refresh message disable extensions, an additional capability bit is added to the common RSVP header, which is defined in [RFC2205].

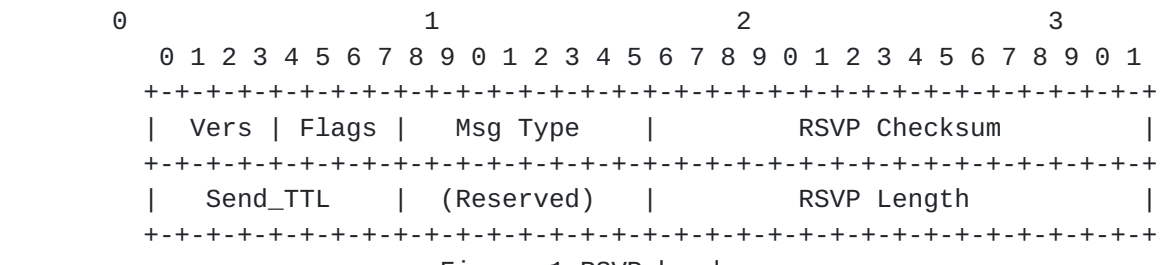

Figure 1 RSVP header

Flags: 4 bits

0x02: Refresh message disable

When set, indicates that this node is willing and capable of supporting the refresh message disable extensions described in this document. This bit is meaningful only between RSVP neighbors.

Nodes supporting the refresh message disable extensions must also take care to recognize when a next hop stops sending RSVP messages with the Refresh-Message-Disable bit set. To cover this case, nodes supporting the refresh message disable extensions MUST examine the flags field of each received RSVP message. If the flag changes from indicating support to indicating non-support then, unless configured otherwise, MUST NOT stop state refreshes to that neighbor.

Note: Refresh messages can only be disabled if the neighbor node must support both reliable RSVP message delivery in [RFC2961] and Hello message extension in [RFC3209] [RFC4558]. When RSVP refresh messages are disabled, both reliable RSVP message delivery in [RFC2961] and Hello message extension in [RFC3209] [RFC4558] must be enabled.

4.1. Reliable RSVP message delivery

Refresh messages are used to both synchronize state between RSVP neighbors and to recover from lost RSVP messages. The reliability and latency problem occurs when a non-refresh RSVP message is lost in transmission.

An implementation that supports the techniques discussed in this document must support the functionality described in [RFC2961] as follows:

- It MUST support reliable delivery of Path/Resv and the corresponding Tear/Err messages (as specified in Section 4 of [RFC2961]).

- It MUST support retransmission of all unacknowledged RSVP-TE messages using exponential backoff (as specified in Section 6 of [RFC2961]).

4.2. Hello Extension for tear message

When RSVP refresh messages are disabled, the time to deallocate resources after a tear message is lost is an issue. To solve this problem, MUST make use of the Hello session based on the Node-ID ([RFC3209][RFC4558]) for detection of RSVP-TE signaling adjacency failures.

An implementation that supports the techniques discussed in this document must support the functionality as follows:

- MUST make use of the Hello session based on the Node-ID ([RFC3209][RFC4558]) for detection of RSVP-TE signaling adjacency failures. A default value of 9 seconds is RECOMMENDED by this document for the configurable node hello interval (as opposed to the default value of 5 milliseconds proposed in Section 5.3 of [RFC3209]). The Hello message format is as follows:

  <Hello Message> ::= <Common Header> [<INTEGRITY>]
  <HELLO>

  The HELLO Object formats is as follows:

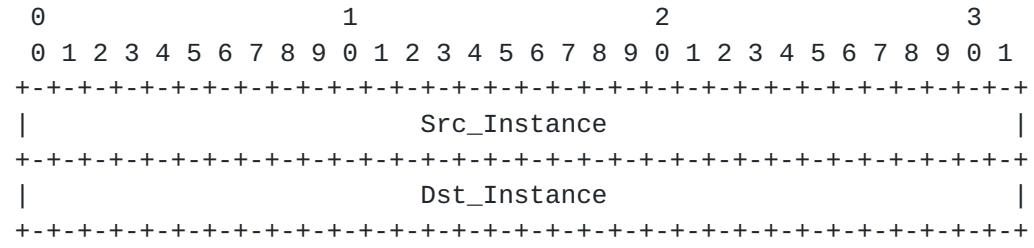

Figure 2 Format for hello object

Src_Instance: 32 bits

a 32 bit value that represents the sender's instance. The advertiser maintains a per neighbor representation/value. This value MUST change when the sender is reset, when the node reboots, or when communication is lost to the neighboring node and otherwise remains the same. This field MUST NOT be set to zero (0).

Dst_Instance: 32 bits

The most recently received Src_Instance value received from the neighbor. This field MUST be set to zero (0) when no value has ever been seen from the neighbor.

- MUST implement coupling the state of individual LSPs with the state of the corresponding RSVP-TE signaling adjacency. When an

RSVP-TE speaker detects RSVP-TE signaling adjacency failure, the speaker MUST act as if all the Path and Resv states learned via the failed signaling adjacency have timed out.

5. Security Considerations

This document does not introduce additional security requirements and mechanisms. Implementation of the mechanism follows the security specification of [RFC2205].

6. IANA Considerations

This document makes no IANA requests.

7. Normative References

[RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <https://www.rfc-editor.org/info/rfc2205>.

[RFC2961] Berger, L., Gan, D., Swallow, G., Pan, P., Tommasi, F., and S. Molendini, "RSVP Refresh Overhead Reduction Extensions", RFC 2961, DOI 10.17487/RFC2961, November 0301, <https://www.rfc-editor.org/info/rfc2961>.

[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, <https://www.rfc-editor.org/info/rfc3209>.

[RFC4558] Ali, Z., Rahman, R., Prairie, D., and D. Papadimitriou, "Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement", RFC 4558, DOI 10.17487/RFC4558, June 2006, <https://www.rfc-editor.org/info/rfc4558>.

[RFC8370] Beeram, V., Ed., Minei, I., Shakir, R., Pacella, D., and T. Saad, "Techniques to Improve the Scalability of RSVP-TE Deployments", RFC 8402, DOI 10.17487/RFC8370, May 2018, <https://www.rfc-editor.org/info/rfc8370>.

Authors' Addresses

Jun Gao
Fiberhome Telecom LTD.
Gaoxin 4th Road 6#
Wuhan, Hubei 430079
China

Email: jgao@fiberhome.com

Jinyou Dai
Fiberhome Telecom LTD.
Gaoxin 4th Road 6#
Wuhan, Hubei 430079
China

Email: djy@fiberhome.com