

teas
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

Q. Xiong
X. Zhu
ZTE Corporation
2 March 2026

RSVP Extensions for Rate-based Resource Quota
draft-xiong-teas-rsvp-resource-quota-00

Abstract

This document proposes RSVP extensions for rate-based resource quota to enable dynamic resource reservation, achieving effective rate control in multi-flow data transmission.

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."

This Internet-Draft will expire on 3 September 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions Used in This Document	3
2.1. Abbreviations and Terminology	3
2.2. Requirements Language	3
3. RSVP Signalling for Rate-based Resource Quota Reservation . .	3
4. RSVP Extensions	5
4.1. Objects Extensions	5
4.1.1. RateControl Object	5
4.1.2. Quota Object	6
5. Security Considerations	7
6. IANA Considerations	8
7. Informative References	8
Authors' Addresses	9

1. Introduction

The increasing number of job-based applications (e.g., AI/HPC, video conferencing, cloud gaming) demand for high-throughput, low-latency communications over WANs. To ensure efficient transmission of these services in the network and avoid congestion, the network can establish tunnels and reserve resources based on the requirements, guaranteeing high-throughput transmission within the job's completion time.

A job decomposed into multiple tasks for parallel transmission results in the concurrent transmission of multiple flows. Traditional congestion control mechanisms implemented on hosts lack rate controls, resulting in unbounded rate adjustments. The static resource allocation methods can provide QoS guarantee to imigrate the congestion for the job. But in multi-flow communications, network bandwidth is shared among multiple flows, leading to suboptimal performance, including slow-flow (tailing) and throughput degradation as follows.

*Slow-flow tailing: when one flow experiences congestion, it can delay others.

*Throughput inefficiency: static bandwidth allocation for multi-flows may over/under-utilize resources.

*Lack of synchronization: multi-flows may arrive with different delays, disrupting the synchronization of group communication.

The network should provide efficient resource reservation mechanisms at nodes along the path. This is not bandwidth allocation but rather quota reservation, which can be used for admission control. Quota is

expressed as a vector of resource quantities (bandwidth, buffer, queue, etc.) for a given time frame. The network can make dynamic resource reservation for different time frames defined by quota. The network nodes can reserve resource quota based on rate control for each flow. For example, it could reserve the minimum bandwidth quota to guarantee the minimum rate of multiple flows and schedule dynamic resources to achieve the maximum rate.

This document proposes RSVP extensions for rate-based resource quota to enable dynamic resource reservation, achieving effective rate control in multi-flow data transmission.

2. Conventions Used in This Document

2.1. Abbreviations and Terminology

This document uses the following terminology:

Resource Quota: A vector of resource quantities (bandwidth, buffer, queue, etc.) for a given time frame.

2.2. 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. RSVP Signalling for Rate-based Resource Quota Reservation

The network needs to perform the admission control based on available capacity after receiving the flow transmission request. The request will be accepted when the rate-based resource quota reservation succeeds; otherwise, it will be rejected. When the network reserves the resource quota based on the minimum rate of the flow and scheduled the resource based on the maximum rate, the traffic can be transmitted within the controlled range of rates to guarantee the completion time and the throughput among multiple flows.

The network performs dynamic resource reservation for different time frames defined by quota as follows.

- * The ingress node sends Path message to request the resource quota while carrying the Quota Object (Section 4.1.1) to indicate the required resource (e.g., bandwidth, buffer, queue) within a time frame. The resource quota can be computed based on the minimum rate of the flow.

- * The transit node receiving a Path message needs to check the resources allocated for the service and transfer it to the next node. When the remaining resources are not enough for the requested quota, it needs to reply with a PathErr message.
- * The egress node receiving the Path message needs to check the output resources and return a Resv message carrying the Quota Object (Section 4.1.1) and RateControl Object (Section 4.1.2) to indicate the maximum rate of the flow.
- * The network nodes receiving the Resv message need to reserve the resource quota within a time frame, but not allocate it while the traffic in the network is still sharing the resources. For example, the nodes need to subtract the resource quota within the time frames and the subtraction result will be viewed as the resource constraints for admission control.
- * The egress node may send a Notify message carrying RateControl Object to update the maximum rate of the flow.
- * When the data transmission is completed, the resource quota will be released with PathTear and ResvTear messages.

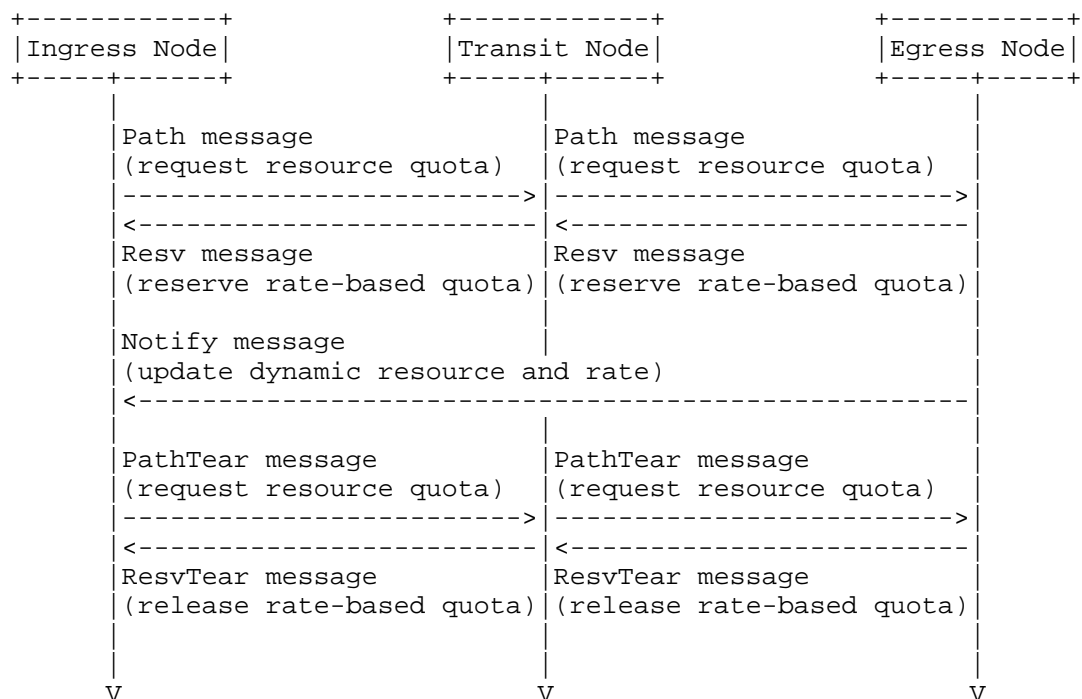


Figure 1 RSVP Signaling for Rate-based Resource Quota Reservation

4. RSVP Extensions

4.1. Objects Extensions

4.1.1. RateControl Object

This document proposes the RateControl Object to carry the rate information which is acknowledged by the egress node. The format of the RateControl Object is shown in Figure 2.




Figure 2 RateControl Object

Min-Rate: 32bits, indicates the minimum rate for the flow.

Max-Rate: 32bits, indicates the maximum rate for the flow.

4.1.2. Quota Object

This document proposes the Quota Object to carry the resource quota information. The format of the Quota Object is shown in Figure 3.

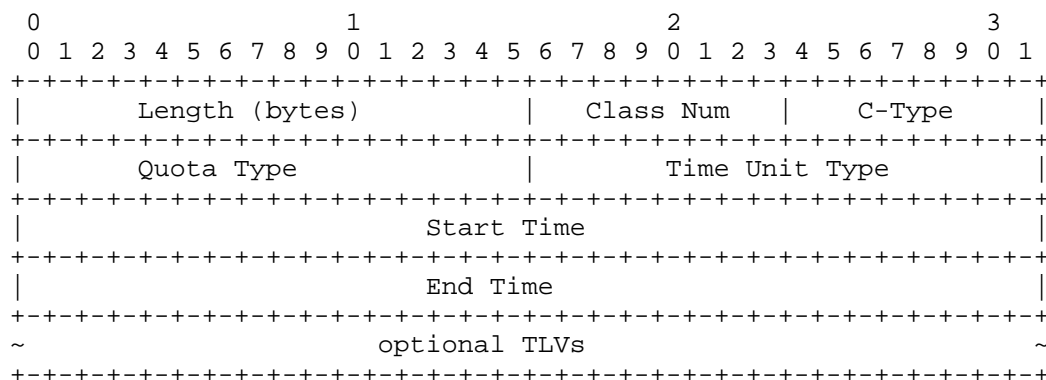


Figure 3 Quota Object

Quota Type: 16 bits, indicates the type of resource quota, including the bandwidth, buffer, queue, CPU size, etc.

Time Unit Type: 16 bits, indicates the type of time unit, including second, microsecond, millisecond, and minute.

Start Time: 32 bits, indicates the start time of the time frame.

End Time: 32 bits, indicates the end time of the time frame.

optional TLVs: variable length and multiple TLVs can be carried to indicate the resource quota.




Figure 4 Bandwidth Quota TLV

Bandwidth: 32 bits, indicates the bandwidth quota which is reserved for the flow.




Figure 5 Time-based Quota TLV

Timeslot/Cycle: 32 bits, indicates the time-based quota which is reserved for the flow.




Figure 6 Queue Quota TLV

Queue ID: 32 bits, indicates the queue quota which is reserved for the flow.




Figure 7 Buffer Quota TLV

Buffer: 32 bits, indicates the buffer quota which is reserved for the flow.




Figure 8 CPU Quota TLV

CPU Size: 32 bits, indicates the CPU quota which is reserved for the flow.

5. Security Considerations

The security considerations specified in [RFC2205] and [RFC4860] apply to this document. In addition, [RFC4230] and [RFC6411] provide useful guidance on RSVP security mechanisms.

The network nodes need to perform quota-based resource computation and reservation based on authentication (e.g., [RFC2747] and [RFC3097]) and authorization (e.g., [RFC6749]).

6. IANA Considerations

TBA.

7. Informative References

- [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/rfc/rfc2119>>.
- [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/rfc/rfc2205>>.
- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, DOI 10.17487/RFC2747, January 2000, <<https://www.rfc-editor.org/rfc/rfc2747>>.
- [RFC3097] Braden, R. and L. Zhang, "RSVP Cryptographic Authentication -- Updated Message Type Value", RFC 3097, DOI 10.17487/RFC3097, April 2001, <<https://www.rfc-editor.org/rfc/rfc3097>>.
- [RFC4230] Tschofenig, H. and R. Graveman, "RSVP Security Properties", RFC 4230, DOI 10.17487/RFC4230, December 2005, <<https://www.rfc-editor.org/rfc/rfc4230>>.
- [RFC4860] Le Faucheur, F., Davie, B., Bose, P., Christou, C., and M. Davenport, "Generic Aggregate Resource ReSerVation Protocol (RSVP) Reservations", RFC 4860, DOI 10.17487/RFC4860, May 2007, <<https://www.rfc-editor.org/rfc/rfc4860>>.
- [RFC6411] Behringer, M., Le Faucheur, F., and B. Weis, "Applicability of Keying Methods for RSVP Security", RFC 6411, DOI 10.17487/RFC6411, October 2011, <<https://www.rfc-editor.org/rfc/rfc6411>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/rfc/rfc6749>>.

[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/rfc/rfc8174>>.

Authors' Addresses

Quan Xiong
ZTE Corporation
Email: xiong.quan@zte.com.cn

Xiangyang Zhu
ZTE Corporation
Email: zhu.xiangyang@zte.com.cn