QUIC Working Group
Internet-Draft
Intended status: Standards Track
Expires: 4 September 2025

Z. Li
China Mobile
W. Cheng
J. Wang
Centec

3 March 2025

QUIC NAT
draft-li-quic-nat-optimization-00

Abstract

QUIC uses UDP as its transport layer protocol. Many existing NAT routers rely on observing TCP SYN, ACK, and FIN packets to determine the establishment and termination of connections, thereby precisely maintaining the lifecycle of NAT mapping entries. For QUIC, NAT devices can only manage mapping entries based on ordinary UDP aging mechanisms, which may cause NAT entries for long-lived QUIC connections to be prematurely aged and re-bound, resulting in changes to source ports. This document proposes a solution that extends the IP header options field to identify QUIC connections, facilitating NAT devices that do not recognize QUIC to accurately determine the lifecycle of QUIC connections and prevent random aging and re-mapping of long-lived QUIC connections.

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 .

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 4 September 2025.

Table of Contents

1. Introduction

QUIC (Quick UDP Internet Connections) is a transport layer protocol that uses UDP and aims to provide reliable transmission similar to TCP. However, because QUIC uses UDP, traditional NAT devices cannot manage connection lifecycles by observing TCP connection establishment and termination signals. This leads to the problem that long-lived QUIC connections passing through NAT devices may be interrupted or experience source port changes due to NAT mapping entries being aged out. Many existing NAT routers rely on observing TCP SYN, ACK, and FIN packets to determine the establishment and termination of connections, thereby precisely maintaining the lifecycle of NAT mapping entries. For QUIC, NAT devices can only manage mapping entries based on ordinary UDP aging mechanisms, which may cause NAT entries for long-lived QUIC connections to be prematurely aged and re-bound, resulting in changes to source ports. This document proposes a comprehensive solution that extends the IP header options field to optimize QUIC NAT traversal. The solution addresses four core issues: 1. Connection Identification: By including a Connection ID, the solution allows NAT devices to accurately determine the lifecycle of QUIC connections, preventing random aging and re-mapping of long-lived QUIC connections. 2. Dynamic Aging Time: An Aging Time field enables the service side to

dynamically maintain session aging times, reducing resource overhead for NAT devices. 3. Service Level Classification: A Service Level field allows for user-defined end-to-end service quality classifications, enabling NAT devices to differentiate service priorities and implement fine-grained quality of service in conjunction with IP DSCP. 4. Security Group Division: A Security Group ID field supports the division of different security groups, allowing network devices to directly implement security policies based on this ID.

2. Problem Statement

Traditional NAT routers rely on observing TCP SYN, ACK, and FIN packets to determine the establishment and termination of TCP connections, thereby precisely maintaining the lifecycle of NAT mapping entries.

For QUIC, which is based on UDP, NAT devices can only manage mapping entries using the standard UDP aging mechanism. UDP aging mechanisms are usually based on fixed timeout intervals and cannot adjust according to the actual state of the connection. This results in the following issues: NAT mapping entries for long-lived QUIC connections may be prematurely aged, requiring re-establishment of the mapping. Re-establishing mapping relationships may lead to changes in source ports, affecting the continuity and performance of QUIC connections. As shown in the figure below (figure to be inserted), the occurrence process of the above problem is depicted.

3. Solution

This proposal extends the IP header options field to mark QUIC connections, allowing NAT devices that do not natively support QUIC to accurately determine the lifecycle of a QUIC connection. This prevents random aging and re-mapping of long-lived QUIC connections by the NAT device.

3.1. Analysis of QUIC NAT Packet Format

The format of a QUIC packet is shown in the figure below. The Connection ID field in the QUIC header is used to identify a QUIC connection. This field has a variable length, which can be 0, 1, 4, or 8 bytes, depending on the encoding of the preceding Flags field.

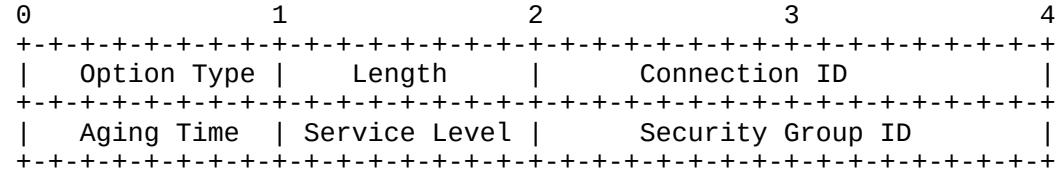

Figure 1: QUIC NAT Packet Format of IP Option

Connection ID, The Connection ID field is used to identify a QUIC connection, allowing NAT devices to accurately recognize and maintain the lifecycle of QUIC connections. This prevents random aging and re-mapping of long-lived QUIC connections by NAT devices that do not natively support QUIC protocol recognition.

Aging Time, The Aging Time field represents the aging time of each session. It is dynamically maintained by the service side, reducing the resource overhead and cost for NAT devices in terms of session state statistics and monitoring.

Service Level, The Service Level field indicates the service quality classification for the connection. This allows users to specify end-to-end service quality classifications. NAT devices can use this field to differentiate service priorities and implement fine-grained quality of service in conjunction with IP DSCP, enabling hierarchical multi-level scheduling.

Security Group ID, The Security Group ID field is used to divide different security groups. Users can specify different security groups, and network devices can directly implement security policies based on the Security Group ID. This reduces the resource cost and complexity of security matching table entries on devices.

4. Conclusion

This document proposes a QUIC NAT optimization solution based on extending the IP header options field. By introducing fields for Connection ID, Aging Time, Service Level, and Security Group ID, this solution addresses issues related to connection lifecycle management, resource optimization, service quality differentiation, and security policy implementation in NAT environments for QUIC connections. This comprehensive approach ensures connection stability, performance, and security while reducing the burden on NAT devices.

5. Security Considerations

TBD.

6. IANA Considerations

TBD.

Authors' Addresses

Zhiqiang Li
China Mobile
Beijing
100053
China
Email: lizhiqiangyjy@chinamobile.com

Wei Cheng
Centec
Suzhou
215000
China
Email: chengw@centec.com

Junjie Wang
Centec
Suzhou
21500
China
Email: wangjj@centec.com