

QUIC Working Group
Internet-Draft
Intended status: Standards Track
Expires: 10 May 2026

Z. Li
Z. Du
China Mobile
W. Cheng
J. Wang
Centec
6 November 2025

CDN Node Selection Enhancement using Anycast and QUIC
draft-li-cdn-node-selection-enhancement-00

Abstract

Content Delivery Networks (CDNs) are critical infrastructure for improving user experience by serving content from geographically proximate servers. Traditional CDN node selection mechanisms, such as DNS-based redirection and BGP-based anycast, primarily rely on network proximity and often lack awareness of server load or specific application requirements. This can lead to suboptimal performance and inefficient resource utilization.

This document proposes an enhanced mechanism for dynamic CDN node selection that leverages IPv4/v6 anycast, BGP, and the QUIC transport protocol. The proposed solution enables CDN nodes to advertise or withdraw their anycast routes based on real-time load and link quality. Furthermore, it defines a mechanism for clients to signal their service requirements (e.g., high-compute, low-latency, or high-bandwidth) and for servers to redirect clients to more optimal nodes using QUIC's connection migration capabilities. This allows for a more intelligent, load-aware, and application-aware CDN node selection process.

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 10 May 2026.

Copyright Notice

Copyright (c) 2025 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. Problem Statement | 3 |
| 3. Solution | 4 |
| 4. Best Effort (BE) Mode | 4 |
| 5. Traffic Engineering (TE) Mode | 5 |
| 6. Security Considerations | 5 |
| 7. IANA Considerations | 5 |
| Authors' Addresses | 5 |

1. Introduction

With the proliferation of latency-sensitive and resource-intensive applications, the performance of application servers has become a decisive factor in user engagement and loyalty. CDNs enhance user experience by caching content closer to end-users, thereby reducing latency. The efficiency of a CDN system heavily relies on its ability to select an optimal server node for each client request.

Current prevalent node selection mechanisms include:

1. DNS-based redirection: Resolves a hostname to the IP address of a topologically "close" server, typically based on the client's DNS resolver location.

2. Anycast routing: A single IP address is assigned to multiple servers. Network routing protocols direct the client's packets to the server that is "closest" in terms of network topology.

These methods, while effective at reducing network latency, generally fail to account for the real-time load on CDN servers or the specific needs of different applications (e.g., compute-intensive vs. bandwidth-intensive).

This document specifies a solution that integrates BGP-controlled anycast with advanced features of the QUIC protocol. It enables load-aware route advertisement and provides a standardized way for clients to be redirected to more suitable CDN nodes without terminating their session, leveraging QUIC's connection migration feature.

2. Problem Statement

DNS-based and anycast-based CDN selection schemes primarily optimize for network distance.

This document aims to address these issues by creating a more dynamic, load-aware, and application-aware selection mechanism that minimizes service disruption.

Existing CDN node selection solutions have the following limitations:

1. *DNS-based Redirection and Basic Anycast*:

- * Select nodes based on network distance (not server load), leading to load imbalance;
- * Cannot differentiate service types (e.g., low-latency vs. high-bandwidth), resulting in suboptimal service quality;
- * Poor real-time performance and scalability;
- * Route reconfiguration often causes connection interruptions.

2. *Service-specific Optimization*:

- * Typically non-standard (e.g., video-only protocols), leading to poor compatibility and universality across different services.

3. Solution

This document proposes a distributed dynamic CDN node selection method based on IPv4/v6 Anycast and QUIC. The method considers server load (CPU utilization, remaining computing power), service requirements (high bandwidth, low latency, high computing power), and supports seamless scaling. Key network elements include:

- * ***CDN Client***: Initiates service requests via IPv4/v6 Anycast and QUIC, processes responses, and supports QUIC connection migration.
- * ***CDN Server***: Configured with IPv4/v6 Anycast addresses; advertises/withdraws Anycast routes via BGP based on load; receives and responds to QUIC requests; caches resources from origin servers.

The solution supports two working modes: BE (Best Effort) and TE (Traffic Engineering).

4. Best Effort (BE) Mode

In BE mode, the system provides basic load-aware routing and resilience.

1. Each CDN Node is configured with an IPv4/v6 anycast address.
2. The CDN Node dynamically manages the BGP advertisement of its anycast address based on its status. For example:
 - * If the CPU load is below a threshold (e.g., 80%) and the link to its gateway router is not congested (e.g., $RTT \leq 1.5 \times$ baseline RTT), the node advertises a route for the anycast address via BGP.
 - * If the CPU load exceeds the threshold or the link becomes congested, the node withdraws the BGP route.
3. A client sends its initial QUIC packets to the anycast address. BGP directs these packets to the topologically nearest CDN Node that is currently advertising the route.
4. If the serving CDN Node becomes unavailable (e.g., because it withdraws its BGP route due to overload), the client's subsequent packets will be routed to the next-nearest available node. The client detects this path change and uses QUIC's connection migration feature RFC9000 to seamlessly continue the session with the new node without re-establishing the connection.

5. Traffic Engineering (TE) Mode

TE mode extends BE mode by allowing for more granular, application-aware node selection.

In TE mode, CDN Nodes are assumed to have knowledge of their ranking relative to other nodes within a specific domain (e.g., all nodes serving a particular region). This ranking is based on three key metrics:

- * Compute Power (Kn): A rank indicating the node's available computing capacity (e.g., based on FLOPS). A lower Kn value might indicate higher available compute.
- * Latency (Dn): A rank based on the baseline network latency from the node to its gateway, indicating its proximity to the network edge. A lower Dn value indicates lower latency.
- * Bandwidth (Bn): A rank representing the node's available network bandwidth. A lower Bn value might indicate higher available bandwidth.

The specific mechanism by which nodes discover each other and determine these ranks (e.g., via a centralized controller or a distributed discovery protocol) is outside the scope of this document.

6. Security Considerations

TBD.

7. IANA Considerations

TBD.

Authors' Addresses

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

Zongpeng Du
China Mobile
Beijing
100053
China
Email: duzongpeng@chinamobile.com

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

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