

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 4 May 2026

W. Yang
Y. Ren
X. Zhou
CNIC
Q. Wu
ICT
G. Xie
CNIC

31 October 2025

A Packet Marking Policy for Low Latency, Low Loss, and Scalable
Throughput (L4S)
draft-yang-l4s-packet-marking-policy-00

Abstract

Low Latency, Low Loss, and Scalable Throughput (L4S)[RFC9330] is a technology designed to mitigate queueing delays while maintaining high throughput for IP traffic. In real-time communication (RTC) applications over 5G networks, rapidly fluctuating wireless link conditions impose strict requirements on congestion control algorithms, which must simultaneously ensure low latency and high bandwidth utilization. This document proposes a packet marking policy for L4S in 5G networks, where intermediate base station devices compute a link load factor and use it as a probabilistic signal to mark packets in L4S flows.

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 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. Terminology	3
3. A New L4S Marking Process	3
3.1. Deployment Challenges in 5G	4
3.2. Link Load Factor Marking Strategy	4
4. IANA Considerations	5
5. Security Considerations	5
6. Acknowledgements	5
7. Normative References	5
Appendix A. Historical Note	6
Authors' Addresses	6

1. Introduction

Low-latency real-time communication (RTC) applications, such as interactive video, cloud gaming, and AR/VR, require efficient congestion control mechanisms that can react quickly to changes in network conditions. Accurate and timely congestion signals from intermediate network elements are essential to prevent buffer build-up and to maintain both low latency and high throughput.

The Low Latency, Low Loss, and Scalable Throughput (L4S) architecture RFC9330 [RFC9331] provides a framework for reducing queuing delays by using Explicit Congestion Notification (ECN) as an end-to-end congestion signal. However, existing L4S marking approaches are primarily based on queuing delay measurements and do not fully capture the dynamics of wireless access networks. In particular, 5G access networks introduce rapid fluctuations in link capacity, where queuing delay alone is not a sufficient indicator of congestion. As a result, conventional L4S marking methods may provide delayed or imprecise feedback, limiting their effectiveness for RTC applications.

This document proposes a packet marking strategy for L4S tailored to 5G environments. In this design, base stations compute a link load factor and use it as a probabilistic signal to mark packets in L4S flows. The link load information is encoded into packets at the 5G base station and conveyed to the sender via ACK-based feedback. On the sender side, a dynamic send rate adaptation algorithm adjusts the transmission rate based on the reported link load factor, preventing network congestion while balancing throughput and latency.

2. Terminology

RTC: real-time communication.

L4S: Low Latency, Low Loss, and Scalable Throughput (L4S) as defined in [RFC9330] and [RFC9331].

ECN: Explicit Congestion Notification, defined in [RFC3168], [RFC9330], and [RFC9331], used to signal congestion without dropping packets.

gNB: Next-generation Node B, the 5G base station.

MAC: Medium Access Control layer in the 5G protocol stack, responsible for multiplexing and scheduling data flows.

RLC: Radio Link Control layer in the 5G protocol stack, responsible for segmentation, retransmission, and queue management.

PDCP: Packet Data Convergence Protocol layer in the 5G protocol stack, which handles header compression, security, and delivery of IP packets. In this document, the PDCP layer is considered for applying ECN markings, as IP headers remain visible before encryption.


ACK: Acknowledgment message, used in feedback to the sender to convey delivery status and congestion information.

The keywords "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] and [RFC8174].

3. A New L4S Marking Process

3.1. Deployment Challenges in 5G

Low Latency, Low Loss, and Scalable Throughput (L4S) is a congestion control framework designed to reduce queuing delay in general-purpose IP networks. It uses Explicit Congestion Notification (ECN) bits in the IP header to provide end-to-end congestion signaling, enabling applications to adjust their sending rates dynamically without incurring excessive delay or packet loss.



Deploying L4S within the 5G protocol stack introduces specific challenges. In a typical 5G setup, L4S may be realized by maintaining queuing at the RLC layer and applying ECN markings at the PDCP layer, with flow identification based on ECN bits in compliance with existing standards. Although the L4S specification does not prescribe a packet marking policy, most current implementations rely on queuing delay measurements to probabilistically mark packets. In 5G networks, the air interface and wired backhaul mismatch often makes the RLC layer a congestion bottleneck. Ideally, marking should occur at the RLC layer to reflect actual buffer conditions. However, since RLC data is encrypted, modifying ECN bits at that layer is infeasible. A practical alternative is to apply marking at the PDCP layer, where the IP header is still accessible before encryption.

3.2. Link Load Factor Marking Strategy

To address the above limitations, this document specifies a probabilistic marking strategy based on a periodically computed `_link` load factor (P Mark):

$$P_Mark = (T_Size + R_Length) / (Bandwidth \times T_Interval)$$

The parameter `P_Mark` represents the ratio between the traffic offered to the wireless link and the link capacity available over a given interval. It indicates how heavily the link is loaded relative to its service capability.

Where:

- * `T_Interval`: The measurement interval.
- * `T_Size`: The amount of new traffic arriving during `T_Interval`.
- * `R_Length`: RLC real-time queue length before `T_Interval`
- * `Bandwidth`: The average wireless bandwidth allocated by the base station to the current user within the `T_Interval`

If $P_Mark > 1$, all packets are marked; if $0 < P_Mark < 1$, packets are marked with probability `P_Mark`. The marking is performed at the PDCP layer, and the L4S marking information is then fed back to the sender through ACK, allowing the sender to implement the corresponding adaptive rate control.

This strategy provides a practical and compatible path for deploying L4S in 5G systems without altering the core structure of existing base stations.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

For further study.

6. Acknowledgements

Thanks to Wenji Du and Baosen Zhao for discussions and comments on the design of this draft.

7. Normative 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/info/rfc2119>>.

- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<https://www.rfc-editor.org/info/rfc3168>>.
- [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/info/rfc8174>>.
- [RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <<https://www.rfc-editor.org/info/rfc9330>>.
- [RFC9331] De Schepper, K. and B. Briscoe, Ed., "The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9331, DOI 10.17487/RFC9331, January 2023, <<https://www.rfc-editor.org/info/rfc9331>>.

Appendix A. Historical Note

Authors' Addresses

Wanghong Yang
CNIC
Beijing
China
Email: yangwanghong@cnic.cn

Yongmao Ren
CNIC
Beijing
China
Email: renyongmao@cstnet.cn

Xu Zhou
CNIC
Beijing
China
Email: zhouxu@cstnet.cn

Qinghua Wu
ICT
Beijing
China
Email: wuqinghua@ict.ac.cn

Gaogang Xie
CNIC
Beijing
China
Email: xie@cnic.cn