

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 23 October 2025

K. Akiyama
R. Shinkuma
HDT/SIT
21 April 2025

Content-based IP-Multicast Grouping Framework for Real-time Spatial
Sensing and Control Applications with Edge Computing
draft-akiyama-cmg-01

Abstract

This document describes a content-based multicast grouping framework aimed at simplifying routing control and reducing unnecessary traffic in large-scale networks for real-time spatial sensing and control applications. The framework introduces content-based multicast groups, which are managed at the local level by routers without the need for global group membership tracking. Additionally, a "topic" concept is introduced, allowing routers to manage multicast delivery based on data content. This framework reduces bandwidth consumption and simplifies multicast routing while offering flexible data delivery across various topics.

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 23 October 2025.

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	2
3. Multicast Framework	3
3.1. Conventional method	3
3.2. Content-based Multicast	3
3.3. Content-based Multicast Grouping	3
4. Example Use Case	4
5. IANA Considerations	5
6. Security Considerations	5
7. Conclusion	5
8. Acknowledgements	6
9. Normative References	6
Authors' Addresses	6

1. Introduction

Multicast is an efficient data transmission method for distributing data to multiple receivers simultaneously. The Internet Protocol (IP) multicast service model is defined in RFC 1112 [RFC1112]. RFC 5110 [RFC5110] describes an overview of multicast routing architectures. However, these traditional multicast frameworks suffer from several challenges, particularly with global address management and excessive traffic due to group membership updates. In large-scale networks, this can lead to inefficiencies in routing and bandwidth consumption. In real-time spatial sensing and control applications, efficient data transmission and processing are crucial for achieving low latency and high accuracy.

This document introduces a new approach where multicast groups are content-based, and defined locally by routers. It eliminates the need for global multicast address management. Additionally, the concept of a "topic" is introduced to manage data delivery based on the content, offering further flexibility and scalability for multicast applications.

2. Terminology

- * Global Multicast Group: A set of hosts interested in receiving data from the same source.

- * Content-based Multicast Group: A multicast group of directly connected hosts that is content-based and defined and managed by a local router.
- * Topic: A label representing the content of the data being multicast. A topic may correspond to specific data types.
- * Membership Request: A message sent by a host to its nearest router to either join or leave a multicast group.
- * RP: Rendezvous point

3. Multicast Framework

3.1. Conventional method

This section describes the conventional method. When a host wants to receive a transmission through multicast group with a certain group address, it indicates its interest to its first-hop router. The router next initiates hop-by-hop forwarding state. The entire distribution for each group is managed by a RP and the transmissions may be optimized later.

3.2. Content-based Multicast

The introduction of the "topic" concept allows routers to manage multicast streams based on the content or purpose of the data, rather than by the group of hosts interested in receiving the data. Each topic is associated with a content-based multicast group, and multiple hosts can subscribe to the same topic.

A router forwards multicast data based on the active topics, and each router in the delivery chain is aware of which topics are active within its local network. Hosts can subscribe to multiple topics, and routers dynamically manage group membership per topic.

3.3. Content-based Multicast Grouping

In the framework, each router is responsible for managing multicast groups for hosts directly connected to it. Unlike traditional multicast, where routers must track global group memberships, routers in this framework only maintain the status of their next-hop hosts. Group addresses can be reused across different local networks, reducing the complexity of address management.

When a host wishes to join a multicast group (associated with a topic), it sends a membership request to its next-hop router. The router finds the requested content-based multicast group. If the

group exists, that is, if it is already receiving the desired the topic data, it adds the host to the content-based multicast group and begins forwarding the data to the host via multicast. Otherwise, the router creates its own new content-based multicast group for the topic, sequentially sends a membership request to its upstream router, and then begins forwarding the data to the new host.

Similarly, when a host wishes to leave a multicast group, it sends a leave request to its next-hop router. The router removes the host from the content-based multicast group. If no other hosts connected directly to the router are subscribed to the same topic, the router removes the group, stops forwarding the data stream, and sends a leave request to its upstream router. The upstream router then checks its own content-based multicast group for the topic. If there are no more local subscribers, it will cease forwarding the data stream for that topic as well. This process ensures that multicast traffic is only delivered when there are active subscribers, preventing unnecessary bandwidth consumption.

4. Example Use Case

This section describes an example use case based on the network topology shown in Figure 1.




Figure 1: Example network topology

The network consists of three routers (R1, R2, and R3), two visual sensors (VS1 and VS2), an edge computer (EC1), and two users (U1 and U2). Visual sensors, such as light detection and ranging (LiDAR) or cameras, provide real-space data for generating a digital twin on EC1. The users utilize the digital twin in real time for various applications, autonomous driving for example. The devices are connected across several subnets: VS1, VS2, and R1 are linked via the 192.168.0.0/24 subnet; R1, R2, and R3 are linked via the 192.168.1.0/24 subnet; EC1 and R2 are linked via the 192.168.2.0/24 subnet; and R3, U1, and U2 are linked via the 192.168.3.0/24 subnet.

EC1 is assumed to receive information from visual sensors to create the digital twin. VS1 sends its acquired data to a multicast group with the address 224.0.0.1 in the 192.168.0.0/24 subnet, to which R1 belongs. R1 forwards the data to a multicast group with the address 224.0.0.2 in the 192.168.1.0/24 subnet, to which R2 belongs.

Subsequently, R2 forwards the data to a multicast group with the address 224.0.0.3 in the 192.168.2.0/24 subnet, allowing EC1 to receive the data.

It is also assumed that EC1 distributes the digital twin data to U1. EC1 sends the digital-twin data to a multicast group with the address 224.0.0.4 in the 192.168.2.0/24 subnet, to which R2 belongs. R2 forwards the data to a multicast group with the address 224.0.0.5 in the 192.168.1.0/24 subnet, to which R3 belongs. Finally, R3 forwards the data to the multicast group with the address 224.0.0.6 in the 192.168.3.0/24 subnet, allowing U1 to receive them.

If U2 wishes to access the digital twin data, it notifies its first-hop router, R3. U2 joins to the multicast group with the address 224.0.0.6 and R3 can receive the data.

If U2 wishes to access visual-sensor data from VS1, it notifies its first-hop router, R3. As R3 does not have the data to forward, it propagates this notification to downstream routers. Then R3 joins to the multicast group with the address 224.0.0.2 and R3 can receive the data. R3 establishes a new multicast group with the address 224.0.0.7 in the 192.168.3.0/24 subnet. Finally, U2 joins to the group and can receive the data.

5. IANA Considerations

This document makes no request of IANA. Note to RFC Editor: this section may be removed on publication as an RFC.

6. Security Considerations

Security in the multicast framework must be addressed to prevent unauthorized access to multicast streams. Routers should implement access control mechanisms to verify whether a host is authorized to join a specific topic. Encryption of multicast traffic should also be considered to prevent eavesdropping on sensitive data.

7. Conclusion

The content-based multicast framework provides a flexible and efficient solution for multicast communication in large-scale networks. By managing multicast group membership locally and introducing the topic-based data delivery model, the framework reduces bandwidth consumption, simplifies routing, and enhances the scalability of multicast services.

8. Acknowledgements

This work is supported in part by the National Institute of Information and Communications Technology, Japan (JPJ012368C06401).

9. Normative References

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.

[RFC5110] Savola, P., "Overview of the Internet Multicast Routing Architecture", RFC 5110, DOI 10.17487/RFC5110, January 2008, <<https://www.rfc-editor.org/info/rfc5110>>.

Authors' Addresses

Kuon Akiyama
Hyper Digital Twins Co., Ltd. / Shibaura Institute of Technology
Toyosu Campus
3-7-5 Toyosu, Koto-ku, Tokyo, Tokyo
135-8548
Japan
Email: k.akiyama@hyper-digitaltwins.com

Ryoichi Shinkuma
Hyper Digital Twins Co., Ltd. / Shibaura Institute of Technology
Toyosu Campus
3-7-5 Toyosu, Koto-ku, Tokyo, Tokyo
135-8548
Japan
Email: shinkuma@hyper-digitaltwins.com