

SPRING
Internet-Draft
Intended status: Standards Track
Expires: 1 September 2026

Z. Li
Z. Du
China Mobile
W. Cheng
J. Wang
G. Zhang
Centec Networks
28 February 2026

SRv6 Extensions for RDMA Multicast Delivery
draft-li-spring-rdma-multicast-over-srv6-00

Abstract

This document specifies SRv6 (Segment Routing over IPv6) extensions for multicast delivery of RDMA (Remote Direct Memory Access) Reliable Connection (RC) traffic. It defines a new SRv6 endpoint behavior, End.MT, that performs per-receiver RDMA Base Transport Header (BTH) modifications at edge nodes of the multicast tree. It also specifies procedures for hop-by-hop aggregation of RDMA ACK, NACK, and CNP response messages along the reverse path. Together, these extensions allow RDMA RC endpoints to communicate using standard point-to-point Queue Pair (QP) semantics while the network distributes data packets over an IP multicast tree. Target deployment scenarios include multi-replica distributed storage writes, HPC collective communications, AI training parameter distribution, and large-scale inference KV cache distribution.

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 1 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	3
1.1. Relationship to Other Work	4
1.2. Requirements Language	4
2. Terminology	4
3. Applicability	5
4. Architecture	6
4.1. Network Roles	6
4.2. Reference Topology	6
5. Data Plane Specification	7
5.1. Multicast Group Setup	7
5.2. Downstream Data Forwarding	7
5.3. Reverse-Path Response Processing	8
5.3.1. ACK Aggregation	9
5.3.2. NACK Aggregation	9
5.3.3. CNP Aggregation	9
6. SRv6 End.MT Behavior	9
6.1. Definition	10
6.2. Pseudocode	10
7. Packet Formats	11
7.1. SRH Usage	11
7.2. End.MT TLV Format	11
8. Intermediate Node State Requirements	12
9. Security Considerations	13
10. IANA Considerations	13
10.1. SRv6 Endpoint Behavior	13
10.2. SRH TLV Type	14
11. References	14
11.1. Normative References	14
11.2. Informative References	15
Acknowledgments	16
Authors' Addresses	16

1. Introduction

Large-scale distributed computing deployments, including data center interconnection, distributed AI training and inference, and national-scale computing networks, rely on high-throughput data transport. RDMA (Remote Direct Memory Access) provides kernel-bypass data transfer with low CPU overhead on both sending and receiving hosts. RDMA one-sided operations, where receive buffers are pre-registered at the receiver's Network Interface Card (NIC), further reduce CPU involvement at the receiving end.

Many distributed applications exhibit one-to-many traffic patterns, including multi-replica storage writes, HPC collective communications (broadcast, scatter), AI training parameter distribution, and KV cache distribution in inference pipelines. IP multicast delivery of such traffic can reduce total network bandwidth consumption compared to per-receiver unicast replication at the source.

The RDMA Reliable Connection (RC) transport mode is the most widely adopted RDMA mode because it supports the complete set of RDMA operations: Read, Write, and Atomic. However, each RC Queue Pair (QP) is a point-to-point association between exactly one sending QP and one receiving QP. RC packets carry per-connection identifiers in the Base Transport Header (BTH), specifically the Destination Queue Pair Number (QPN) and Packet Sequence Number (PSN). These per-connection fields prevent direct application of IP multicast replication to RC traffic because each receiver requires its own QPN and independently tracks PSN state.

Existing application-layer approaches in distributed frameworks (MPI, NCCL, Spark) address this limitation in two ways: by opening separate RC QP connections to each receiver, which results in source bandwidth consumption proportional to the number of receivers, or by constructing application-layer relay trees (tree or ring topologies), which introduce per-hop host-stack traversal latency and additional memory copy overhead at relay nodes.

This document specifies SRv6 extensions that bridge the gap between RDMA RC point-to-point semantics and IP multicast one-to-many delivery. Edge nodes of the multicast tree execute a new SRv6 endpoint behavior (End.MT) that rewrites per-receiver RDMA BTH fields in replicated packet copies. Intermediate and edge nodes aggregate reverse-path RDMA response messages (ACK, NACK, CNP) before they reach the source. RDMA RC endpoints are not required to implement any multicast-specific extensions.

1.1. Relationship to Other Work

The Segment Routing Replication segment defined in [RFC9524] provides a general-purpose SRv6 packet replication behavior (End.Replicate). The End.MT behavior specified in this document is complementary: it performs RDMA-specific BTH header modifications in addition to packet replication at edge nodes. Transit nodes in the multicast tree MAY use End.Replicate or any other IP multicast forwarding mechanism for tree-interior replication.

Fast Congestion Notification Packet (Fast CNP) mechanisms for RoCEv2 networks, such as those described in [I-D.xiao-rtgwg-rocev2-fast-cnp], define switch-originated CNPs sent directly to the sender on a point-to-point basis. The reverse-path CNP aggregation specified in this document operates on the multicast tree topology and is independent of, and compatible with, Fast CNP on individual links.

RoCEv2-based collective communication offloading, as described in [I-D.liu-nfsv4-rocev2], implements in-network aggregation functions for collective operations. This document differs in scope: it addresses one-to-many data distribution (multicast) rather than many-to-one aggregation (reduce), and it does not require RDMA connections between hosts and switches.

1.2. 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.

2. Terminology

This document uses the following terms. Familiarity with SRv6 terminology from [RFC8402], [RFC8754], and [RFC8986] is assumed.

RDMA: Remote Direct Memory Access, as specified in the InfiniBand Architecture [ROCEV2].

RoCEv2: RDMA over Converged Ethernet version 2, an RDMA transport encapsulated in UDP/IPv6 (or UDP/IPv4).

RC: Reliable Connection, an RDMA transport mode providing connection-oriented reliable delivery.

QP: Queue Pair, the RDMA communication endpoint consisting of a Send

Queue (SQ) and a Receive Queue (RQ).

QPN: Queue Pair Number, a 24-bit identifier for a QP.

BTH: Base Transport Header, the RDMA transport header containing the opcode, Destination QPN, PSN, and other fields.

AETH: ACK Extended Transport Header, an RDMA header carrying acknowledgment information.

PSN: Packet Sequence Number, a 24-bit sequence number used for ordering and acknowledgment in RDMA reliable transport.

ACK: Acknowledgment, an RDMA response confirming successful reception.

NACK: Negative Acknowledgment, an RDMA response requesting retransmission.

CNP: Congestion Notification Packet, used in the RoCEv2 ECN-based congestion control mechanism.

SRv6: Segment Routing over IPv6 [RFC8402].

SRH: Segment Routing Header [RFC8754].

SID: Segment Identifier, a 128-bit IPv6 address in SRv6.

End.MT: A new SRv6 endpoint behavior defined in this document for RDMA multicast header transformation at edge nodes.

Designated QPN: A QPN value agreed upon by all multicast group participants during group setup, used as the Destination QPN in data packets traversing the multicast tree.

Proxy Address: An IPv6 address used as the common destination by the source and all receivers to represent the multicast group at the RDMA layer.

3. Applicability

The extensions specified in this document apply to networks that meet all of the following conditions:

1. RDMA transport uses RoCEv2 in Reliable Connection (RC) mode over IPv6.

2. The network underlay supports SRv6 as defined in [RFC8754] and [RFC8986].
3. IP multicast forwarding is available in the underlay between the source and the edge nodes, using any combination of PIM, static multicast routing, or SRv6 Replication segments [RFC9524].
4. Edge nodes are capable of maintaining per-receiver RDMA connection state and performing BTH field modification at line rate.

These extensions do not modify RDMA endpoint behavior. Hosts run unmodified RoCEv2 protocol stacks and establish standard RC QP connections. All multicast-related packet transformations occur within the network.

4. Architecture

4.1. Network Roles

This specification defines the following network roles:

Multicast Source (S): The RDMA sending host. S establishes a standard RDMA RC QP to the Proxy Address using the Designated QPN. S obtains the unicast IPv6 addresses and QPNs of all receivers via a control plane (the control plane protocol is out of scope).

Multicast Receivers (R1..Rn): RDMA receiving hosts. Each Ri establishes a standard RDMA RC QP to the Proxy Address using its own locally assigned QPN.

Edge Nodes: SRv6-capable network nodes adjacent to receivers (e.g., N1-N3 in Figure 1). Edge nodes instantiate the End.MT SID and perform: (a) RDMA BTH Destination QPN replacement, (b) IPv6 Destination Address replacement, (c) ICRC recomputation, and (d) reverse-path ACK/NACK/CNP aggregation.

Transit Nodes: Intermediate forwarding nodes in the multicast tree (e.g., N4-N6 in Figure 1). Transit nodes replicate and forward packets using IP multicast procedures and participate in reverse-path response aggregation.

4.2. Reference Topology




Figure 1: Reference Network Topology

In Figure 1, S1 is the multicast source. R1 through R5 are multicast receivers. N1, N2, and N3 are edge nodes executing the End.MT behavior. N4, N5, and N6 are transit nodes performing IP multicast replication.

5. Data Plane Specification

5.1. Multicast Group Setup

Prior to data transmission, a multicast group MUST be established as follows:

1. All participants (source S1 and receivers R1..Rn) MUST each create an RDMA RC QP directed at the Proxy Address. All QPs MUST use the Designated QPN as the remote (destination) QPN.
2. The source S1 MUST obtain the unicast IPv6 address and the actual local QPN of each receiver Ri via the control plane. The control plane protocol and its signaling procedures are outside the scope of this document.
3. Each edge node MUST be configured to receive IP multicast traffic addressed to the Proxy Address.
4. S1 MUST encode each edge node's associated receiver information (unicast IPv6 addresses and QPNs) into End.MT TLVs within the SRH of data packets, as specified in Section 7.2.

5.2. Downstream Data Forwarding

The source-to-receiver data path operates as follows:

1. S1 constructs RDMA RC data packets with the IPv6 Destination Address set to the Proxy Address and the BTH Destination QPN set to the Designated QPN. S1 encapsulates the packets in an outer IPv6 header with an SRH containing the End.MT SID(s) and associated TLVs.
2. Transit nodes forward the traffic according to their IP multicast forwarding tables, performing tree replication as needed.
3. When a packet arrives at an edge node whose local End.MT SID matches the IPv6 Destination Address, the edge node MUST execute the End.MT behavior (Section 6):
 - a. Parse the End.MT TLV from the SRH to obtain the list of downstream receivers (unicast addresses and QPNs).
 - b. Create one copy of the inner packet for each downstream receiver.
 - c. In each copy, replace the IPv6 Destination Address with the receiver's unicast address.
 - d. In each copy, replace the BTH Destination QPN with the receiver's actual QPN.
 - e. Recompute the Invariant CRC (ICRC) and any other affected checksums.
 - f. Forward each modified copy toward its destination.
4. Each receiver R_i receives a standard RDMA RC unicast packet addressed to its own IPv6 address and QPN. No multicast-specific behavior is required at R_i .

5.3. Reverse-Path Response Processing

Receivers generate three types of response messages toward the source: ACK (acknowledgment of successful reception), NACK (request for retransmission), and CNP (congestion notification). These responses MUST be aggregated hop-by-hop at intermediate nodes before reaching the source, so that the source's retransmission and rate control logic operates correctly without multicast-specific modifications.

5.3.1. ACK Aggregation

The source MUST receive an AckPSN value satisfying the following invariant: for every receiver R_i and every packet with PSN less than or equal to AckPSN, R_i has confirmed successful reception.

Each intermediate node (edge or transit) MUST maintain a record of the most recent AckPSN reported by each downstream branch. When an ACK is received from a downstream branch, the node MUST update the stored AckPSN for that branch. The node MUST forward an ACK upstream carrying an AckPSN equal to the minimum of all downstream branches' stored AckPSN values. The node adjacent to the source (N_6 in Figure 1) MUST write this minimum value into the AETH AckPSN field of the ACK forwarded to the source.

5.3.2. NACK Aggregation

When a receiver detects missing packets, it sends a NACK containing an expected PSN (ePSN) indicating the start of the retransmission range. The source MUST receive an ePSN satisfying the following invariant: for every receiver R_i , all packets with PSN less than ePSN have been successfully received by R_i .

Each intermediate node MUST maintain a per-branch record of ePSN values. For branches that have sent only ACKs (no NACK), the effective ePSN SHOULD be treated as AckPSN + 1. The NACK forwarded upstream MUST carry the minimum ePSN across all downstream branches.

5.3.3. CNP Aggregation

Each intermediate node MUST maintain a per-branch counter (CCount) that records the number of CNP messages received from each downstream branch within a configurable time window T .

At the expiration of each time window T , the node MUST select the branch with the highest CCount value and forward a single CNP upstream representing the most congested downstream path. All CCount values MUST be reset to zero at the start of each new time window.

The time window T MAY be adjusted dynamically based on observed network conditions. The node adjacent to the source MUST rewrite CNP packet headers so that the source processes the CNP as a standard RoCEv2 congestion notification.

6. SRv6 End.MT Behavior

6.1. Definition

End.MT is an SRv6 endpoint behavior instantiated at edge nodes of the RDMA multicast tree. When a node N receives a packet whose IPv6 Destination Address matches a locally instantiated End.MT SID, N performs the processing described in Section 6.2.

End.MT combines the following operations: SRH segment processing, End.MT TLV parsing, per-receiver packet replication, BTH Destination QPN replacement, IPv6 Destination Address replacement, and ICRC recomputation.

6.2. Pseudocode

The following pseudocode follows the conventions of [RFC8986] Section 4.

When N receives a packet destined to S, where S is a local End.MT SID, N does:

```
S01. If NH=SRH and SL > 0 {
S02.   Decrement SL
S03.   Update the IPv6 DA with SRH[SL]           ;; Ref1
S04.   Parse the End.MT TLV associated with S
S05.   Let RecvList = list of (IPv6_Addr, QPN) from TLV
S06.   For each entry (Addr_i, QPN_i) in RecvList {
S07.     Copy the packet                         ;; Ref2
S08.     In the copy, set IPv6 DA = Addr_i
S09.     In the copy, set BTH.DestQPN = QPN_i
S10.     Recompute ICRC over the modified headers
S11.     Forward the copy based on Addr_i        ;; Ref3
S12.   }
S13. } Else {
S14.   Drop the packet                           ;; Ref4
S15. }
```

Ref1: Standard SRH processing per RFC 8754 Section 4.3.1.1.

Ref2: The copy includes all payload beyond the outer IPv6 and SRH headers that are relevant to the inner RDMA frame.

Ref3: FIB lookup on Addr_i determines the outgoing interface.

Ref4: A packet arriving with SL=0 or without SRH is not valid for End.MT processing.

7. Packet Formats

7.1. SRH Usage

This specification uses the standard IPv6 Segment Routing Header (SRH) as defined in [RFC8754]. The SRH carries the End.MT SID in its Segment List and the End.MT TLV in its Optional TLV field.

7.2. End.MT TLV Format

The End.MT TLV is carried in the Optional TLV field of the SRH and conveys per-edge-node receiver information. Its format is as follows:

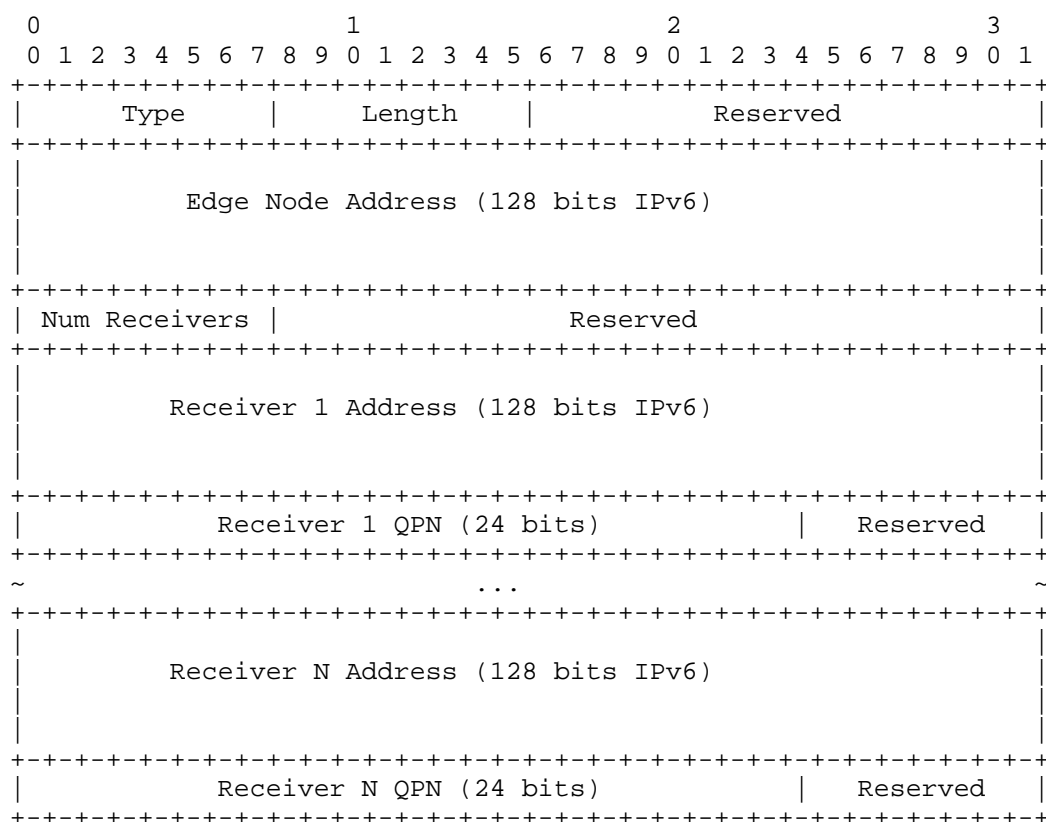


Figure 2: End.MT TLV Format

The fields are defined as follows:

Type (8 bits): SRH TLV type code, to be assigned by IANA (see

Section 10).

Length (8 bits): Length of the Value field in octets, not including the Type and Length fields.

Edge Node Address (128 bits): The IPv6 unicast address of the edge node associated with this TLV entry.

Num Receivers (8 bits): The number of receiver entries following this field.

Receiver i Address (128 bits): The IPv6 unicast address of the i-th receiver.

Receiver i QPN (24 bits): The RDMA Queue Pair Number of the i-th receiver.

Reserved: MUST be set to zero on transmission and MUST be ignored on reception.

Multiple End.MT TLVs MAY be present in a single SRH, one per edge node in the multicast tree. Each End.MT TLV is associated with the End.MT SID of the corresponding edge node.

8. Intermediate Node State Requirements

Each intermediate node (both edge and transit) participating in reverse-path response aggregation MUST maintain the following per-multicast-group state:

Per-branch AckPSN: The most recent AckPSN value received from each downstream branch. Initial value: 0.

Per-branch ePSN: The most recent expected PSN from any NACK received from each downstream branch. For branches that have not sent a NACK, this value SHOULD be set to AckPSN + 1.

Per-branch CCount: A counter of CNP messages received from each downstream branch within the current time window T. Reset to zero at the start of each new time window.

The amount of state is proportional to the number of downstream branches at each node, not to the total number of receivers in the multicast group. Edge nodes additionally maintain the receiver information (addresses and QPNs) learned from the End.MT TLV or from control-plane provisioning.

9. Security Considerations

The security considerations of [RFC8754] and [RFC8986] apply to all SRv6 aspects of this specification. The following additional considerations are specific to RDMA multicast delivery.

The End.MT TLV carries receiver IPv6 addresses and QPNs in the SRH. An on-path attacker able to read SRH contents can obtain receiver topology and RDMA connection identifiers. Implementations operating outside a single administrative trust domain SHOULD protect SRH integrity and confidentiality using the HMAC TLV defined in Section 7 of [RFC8754] or IPsec Encapsulating Security Payload (ESP) encapsulation.

Intermediate nodes maintain per-branch ACK/NACK/CNP aggregation state. An attacker injecting forged response messages could corrupt this state, causing the source to prematurely consider data as acknowledged (via inflated AckPSN) or to trigger unnecessary retransmissions (via forged NACKs). Nodes SHOULD validate that reverse-path response messages originate from addresses within the expected downstream receiver set. BCP 38 [RFC2827] ingress filtering SHOULD be applied at network boundaries.

An attacker injecting a high volume of forged CNP messages could force the source into continuous rate reduction, creating a denial-of-service condition. Intermediate nodes SHOULD implement per-branch CNP rate limiting. The configurable time window T for CNP aggregation provides an inherent dampening effect.

If End.MT TLV contents are modified in transit, packets could be delivered to incorrect RDMA QPs, resulting in data corruption or information disclosure at unintended receivers. The SRH HMAC TLV [RFC8754] provides integrity protection for this purpose. Edge nodes SHOULD verify HMAC before processing End.MT TLVs when operating across trust domain boundaries.

10. IANA Considerations

10.1. SRv6 Endpoint Behavior

This document requests IANA to allocate a new codepoint in the "SRv6 Endpoint Behaviors" sub-registry under the "Segment Routing" registry group [RFC8986]:

Value	Hex	Endpoint Behavior	Reference	Change Controller
TBD1	TBD1	End.MT	[This document]	IETF

Table 1: SRv6 Endpoint Behavior Registration

10.2. SRH TLV Type

This document requests IANA to allocate a new Type value in the "Segment Routing Header TLVs" registry [RFC8754]:

Value	Description	Reference
TBD2	End.MT TLV	[This document]

Table 2: SRH TLV Type Registration

11. References

11.1. 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>>.
- [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>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9524] Voyer, D., Ed., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "Segment Routing Replication for Multipoint Service Delivery", RFC 9524, DOI 10.17487/RFC9524, 28 February 2026, <<https://www.rfc-editor.org/info/rfc9524>>.

11.2. Informative References

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.
- [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>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.
- [I-D.xiao-rtgwg-rocev2-fast-cnp]
Min, X. and H. Li, "Fast Congestion Notification Packet (CNP) in RoCEv2 Networks", Work in Progress, Internet-Draft, draft-xiao-rtgwg-rocev2-fast-cnp-04, 28 February 2026, <<https://datatracker.ietf.org/doc/html/draft-xiao-rtgwg-rocev2-fast-cnp-04>>.
- [I-D.liu-nfsv4-rocev2]
Liu, Y., "RoCEv2-based Collective Communication Offloading", Work in Progress, Internet-Draft, draft-liu-nfsv4-rocev2-00, 28 February 2026, <<https://datatracker.ietf.org/doc/html/draft-liu-nfsv4-rocev2-00>>.

[I-D.hu-rtgwg-rocev2-fcn]

Hu, Z. and Y. Zhu, "Fast Congestion Notification for Distributed RoCEv2 Network Based on SRv6", Work in Progress, Internet-Draft, draft-hu-rtgwg-rocev2-fcn-00, 28 February 2026, <<https://datatracker.ietf.org/doc/html/draft-hu-rtgwg-rocev2-fcn-00>>.

[ROCEV2]

InfiniBand Trade Association, "Supplement to InfiniBand Architecture Specification Volume 1 Release 1.2.1 - Annex A17: RoCEv2", 28 February 2026.

Acknowledgments

The authors thank the members of the SPRING and RTGWG working groups for their review and feedback.

Authors' Addresses

Zhiqiang Li
China Mobile
32 Xuanwumen West Street
Beijing
100053
China
Email: lizhiqiangyjy@chinamobile.com

Zongpeng Du
China Mobile
32 Xuanwumen West Street
Beijing
100053
China
Email: duzongpeng@chinamobile.com

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

Junjie Wang
Centec Networks
Suzhou
215000
China
Email: wangjj@centec.com

Guoying Zhang
Centec Networks
Suzhou
215000
China
Email: zhanggy@centec.com