

Workgroup: BIER Working Group
Internet-Draft:
draft-lamparter-bier-manet-multicast-links
Published: 24 October 2022
Intended Status: Standards Track
Expires: 27 April 2023
Authors: D. Lamparter
in personal capacity (NetDEF, Inc.)
Use, problems and gap analysis for BIER link-layer multicast

Abstract

Mesh networks present a particularly difficult base to provide a multicast service on top of. Not only do normal assumptions of transitive and reflexive connectivity not hold, but proper use of multicast capabilities on lower radio layers can be imperative.

This document provides use case, problem statement and gap analysis for utilizing link-layer multicast features in a BIER domain.

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 27 April 2023.

Copyright Notice

Copyright (c) 2022 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 and use case](#)
- [2. Problem statement](#)
 - [2.1. Reducing BFR-NBR count](#)
 - [2.2. Transmitting to multiple BFR-NBRs](#)
- [3. Gap analysis](#)
- [4. Acknowledgements](#)
- [5. Change Log](#)
- [6. References](#)
 - [6.1. TBD: Normative References](#)
 - [6.2. TBD: Informative References](#)
- [Author's Address](#)

1. Introduction and use case

While the vast majority of internet core networks runs on pointpoint links where link-layer unicast and multicast are essentially identical, radio links present a different environment both in efficiency of medium utilization as well as power use by both transmitters and receivers. Any two receivers reachable by the same transmitter could theoretically also be reached by a single transmission addressed to both of them. At the same time, reachability may be transient on short and long time scales, and is frequently non-transitive or even unidirectional.

The specific characteristics of a radio link ultimately depend on what the link layer technology exposes to the upper layer. This results in huge differences ranging from poorly optimized, costly multicast (e.g. 802.11) to equal cost for unicast or multicast transmission of the same packet (e.g. some satellite networks). Radio links may also have widely divergent numbers of reachable receivers per transmitter. If multicast is to be implemented in these conditions, explicit per-neighbor replication of multicast traffic as described in normal BIER operation can become prohibitively expensive.

As BIER provides a very efficient multicast forwarding system, and may already be in use on non-radio parts of a network, extending BIER for use on radio links is desirable. In particular, even on radio-carried parts of the network, it may be useful to have both unicast and multicast transport options to dynamically switch based on targeted neighbor/receiver counts.

Lastly, the applicable scenarios here can be divided into two broad categories: radio links providing (or emulating) a broadcast domain

with transitive reachability, and radio links omitting this function (general mesh network or NBMA case). The former is a subset of the latter, any solution for mesh/NBMA networks would also work for broadcast radio links, but some aspects could be simplified away.

2. Problem statement

To attempt an efficient solution for the above scenarios, the problem to consider starts out as an abstract quest to reduce radio link use caused by BIER utilizing explicit unicast replication to a number of neighbors (BFR-NBRs) on egress. There are (at least) two angles this can be viewed from:

attempt to reduce the number of BFR-NBRs (BIER neighbors)

attempt to forward data traffic to multiple BFR-NBRs at the same time

2.1. Reducing BFR-NBR count

While sounding odd from a superficial glance, reducing the number of neighbors as seen from BIER protocol operation may be a viable and very simple approach. In particular for broadcast radio links, the entire link could be fused into one BFR-NBR if all members of the link can agree on disjoint bit subsets of forwarding responsibility. Traffic is then sent out as plain broadcast onto the radio link with each receiver masking the bit string to their subset before further processing (possibly discarding the packet entirely if the mask becomes zero.)

TBD: this really doesn't work on non-transitive (i.e. mesh) networks. Explain why.

2.2. Transmitting to multiple BFR-NBRs

The more generic angle is to add some capability for BIER to replace multiple transmissions to distinct BFR-NBRs with one transmission that can reach the same set of BFR-NBRs. Note that not only is this a per-packet consideration, but also not all BFR-NBRs resulting from the BIER forwarding procedure need be reached with the same transmission. In fact, so long as the respective bit strings are disjoint and sum up to the intended set, any number of unicast and multicast transmissions might be combined to implement the forwarding action for a given packet.

Since the decision on which BFR-NBRs shall be used to forward a given packet is fundamentally made by the sender in BIER (as opposed to RPF and receiver-join based approaches), the primary difficulty at this point becomes for the receiver to distinguish whether it was actually an intended recipient of a link-layer multicast packet. The link-

layer unicast destination address would normally perform this distinction, and is now gone. Further, since all recipients would now receive the same bitstring to operate on, the recipient needs to know which bits, if any, it was the intended recipient for.

A secondary difficulty also exists in meeting the requirement for multiple recipients on the same radio link to accept the same encoding for BIER packets. This is more of an encapsulation detail, but generally the receiver is the entity allocating a set of labels which map to SD/BSL/SI. Only if recipients can agree to accept the same encapsulation, multicast delivery becomes possible at all. This problem has been approached with upstream-allocated labels in other context (mLDP), but this is not necessarily the only possible approach.

3. Gap analysis

TBD. Some factors to consider here:

likely only a subset of BIER transports is relevant/useful to specify for MANET operation. For gap analysis, word out some considerations in making those choices later.

MANET routers are generally software forwarders, considerations?

The whole upstream allocated labels thing. But it's not really upstream allocation that's needed, just agreement between downstream receivers. Could be met by just removing configuration variables, e.g. with BIER in IPv6 choose a "global common ground"?

Is a handshake mechanism needed to go in/out of "radio mode"?

If there are switchovers between "plain" and "radio" BIER, will that cause some data packet loss/duplication?

ROLL isn't considered yet at all.

4. Acknowledgements

This document is rooted in discussions with Tony Przygienda and Ronald in 't Velt.

5. Change Log

October 2022 [-00]: Initial text after IETF 114 discussions

6. References

6.1. TBD: Normative References

6.2. TBD: Informative References

Author's Address

David 'equinox' Lamparter
in personal capacity (NetDEF, Inc.)
Leipzig
Germany

Email: equinox@diac24.net