



Tutorial on Bit Index Explicit Replication (BIER)

Daniel Merling, Steffen Lindner, Michael Menth

<http://kn.inf.uni-tuebingen.de>




1. Fundamentals of IP multicast
2. Introduction to BIER
3. BIER Fast Reroute
4. P4 Implementation
5. Summary



► IP multicast (IPMC) efficiently distributes one-to-many traffic

- Forwarding on tree-like structures
 - Individual distribution trees for different IPMC groups
- Only one packet copy per involved link is sent
- IPMC packet does not carry information about receivers



⇒ Traditional IPMC core network requires

1. State per multicast group to know next-hops (NHs) of a packet
2. Signaling in core network when Group subscriptions change

⇒ Scalability of traditional IPMC is limited



- ▶ Bit Index Explicit Replication (BIER) is an efficient transport mechanism for IPMC traffic
 - Under development by the IETF
 - 5 RFCs, 27 active drafts and working group documents

- ▶ Introduces BIER domain
 - Core routers do not require state per IPMC group
 - Only ingress routers participate in signaling when subscribers change

- ▶ Legacy devices do not support BIER
 - SDN allows to implement novel forwarding mechanisms
 - We implemented BIER and BIER fast reroute (BIER-FRR) in P4
 - Software-based for bmv2
 - Hardware-based for Tofino



▶ Bit-forwarding ingress routers (BFIRs)

- Add BIER header to IPMC packet
 - Contains all destinations
 - Individual BIER header per IPMC group
- Receive signals about subscriber changes

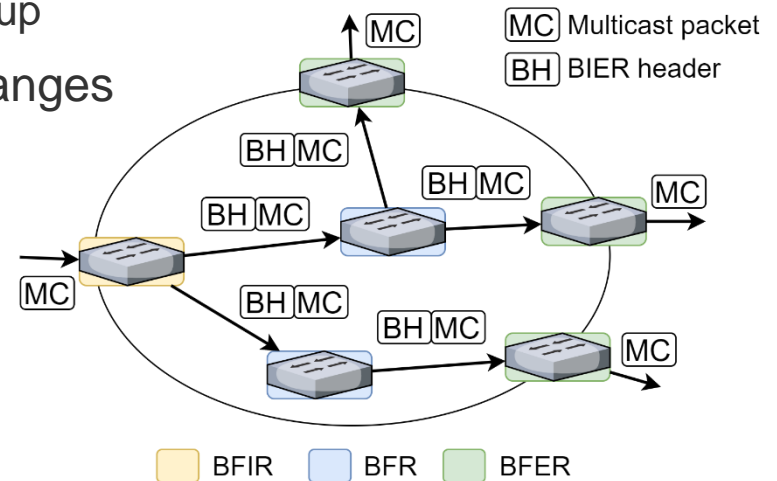
▶ Bit-forwarding routers (BFRs)

- Forward packets to all destinations
- Bit Index Forwarding Table (BIFT)
- Entries based on paths from IGP (routing underlay)

⇒ Forwarding without IPMC-group dependent state

▶ Bit-forwarding egress routers (BFERs)

- Remove BIER header

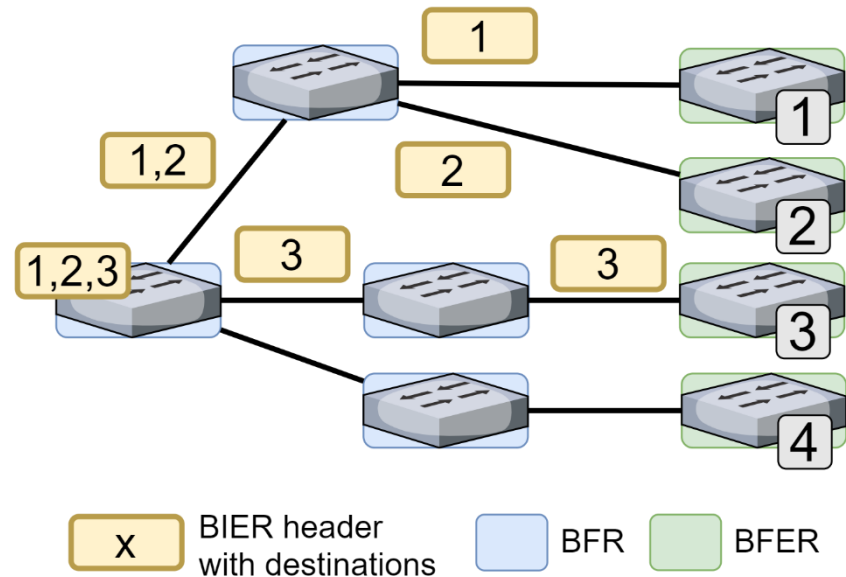




- ▶ BFRs process BIER packets as follows
 - A copy of the BIER packets is forwarded to each NH over which at least one destination of the packet is reached
 - Before a packet copy is forwarded to a NH, all destinations that are reached via other NHs are removed from the packet header (prevents duplicates at receivers)

- ▶ BFRs forward BIER packets to NHs directly on L2
 - No header of the routing underlay is required

- ▶ BIER RFC specifies processing loop
 - Challenging implementation for pipeline-based targets






Problem Statement: Network Failures

► BFRs cannot reach their neighbors

- Entire subtree is disconnected
- Only after BIFT has been recomputed, BFRs can forward BIER traffic again
- Possible only after reconvergence of routing underlay



► Routing underlay restores connectivity faster


- Fast reroute (FRR)
- Faster reconvergence

► BIER fast reroute (BIER-FRR) leverages routing underlay to protect BIER traffic

- Packet is tunneled through routing underlay to downstream nodes in the distribution tree
 - Benefits from FRR or faster reconvergence in routing underlay



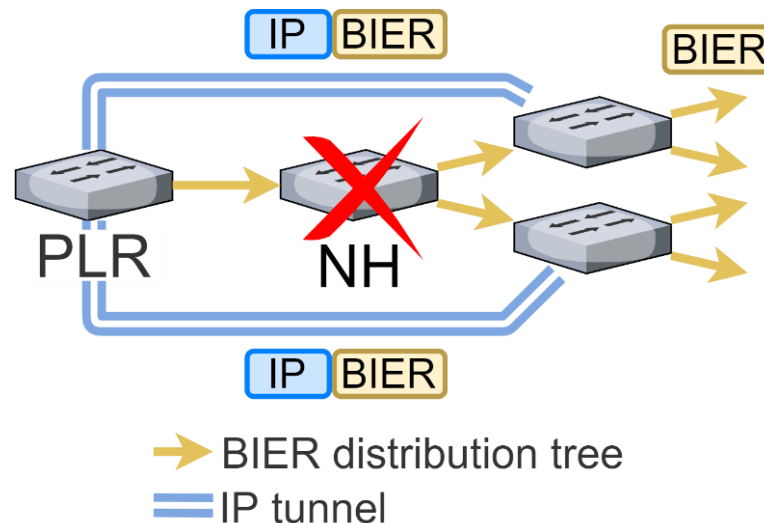
- ▶ Simple mode that protects only against single link failures
- ▶ BFR becomes point of local repair (PLR) when it cannot reach NH
 - PLR tunnels BIER packets through routing underlay to NH by adding header of routing underlay to BIER packet
 - Packet is processed as unicast packet
 - Routing underlay may be affected by failure, too
 - FRR or faster reconvergence restore connectivity quickly
 - NH removes tunnel header and continues BIER forwarding





BIER-FRR with Node Protection (1/2)

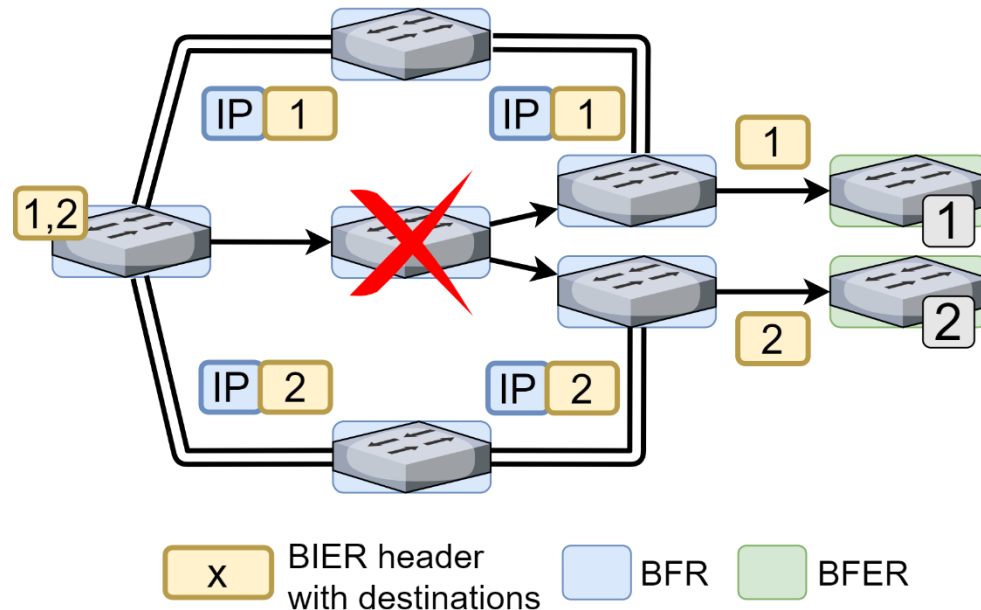
- ▶ Node protection for BIER traffic is more complex than link protection
 - If NH fails, affected subtree cannot be reconnected at NH
 - Next reachable downstream nodes in the distribution tree are next-next-hops (NNHs)
- ▶ BIER packets are tunneled through the routing underlay to all NNHs
 - which should receive a copy of the BIER packet
 - which are on the path towards a destination of the BIER packet





BIER-FRR with Node Protection (2/2)

- ▶ However: failed NH would have adapted BIER header of BIER packets before forwarding
 - PLR bypasses failed NH through a tunnel
 - ⇒ PLR needs to perform additional header adaptations
 - Before tunneling a BIER packet to a NNH, destinations which are reached through other NNHs are cleared from the header





- ▶ Implementation of loop-based BIER processing is challenging for pipeline-based targets
 - Packet recirculation to process packet again
 - Proof-of-concept software implementation for bmv2
 - No quantitative results
 - Qualitative evaluation shows benefit of BIER-FRR in comparison to BIFT recomputation
 - Evaluation for hardware target showed that packet recirculation occupies significant amount of bandwidth
- ⇒ Line rate with workaround based on dedicated packet recirculation ports
 - Ports in loopback mode which use port capacity to process recirculated traffic
 - Number of required recirculation ports depend on amount of multicast traffic



- ▶ BIER is an efficient transport mechanism for IPMC traffic
 - Does not require IPMC-group dependent state in core network

- ▶ In case of failures, updating the BIFT takes much time
 - Routing underlay restores connectivity faster

- ▶ BIER-FRR protects BIER traffic against link and node failures
 - BIER packets are tunneled through routing underlay to N(N)Hs
 - Leverages resilience or faster reconvergence in routing underlay

- ▶ Hardware-based P4 Implementation available
 - BIER in line rate
 - Quick response time in case of failures