

BIER-TE Forwarding

IETF 100, Singapore

Yongqing Zhu, China Telecom
Huannan Chen, China Telecom
Quan Xiong, ZTE
Fangwei Hu, ZTE



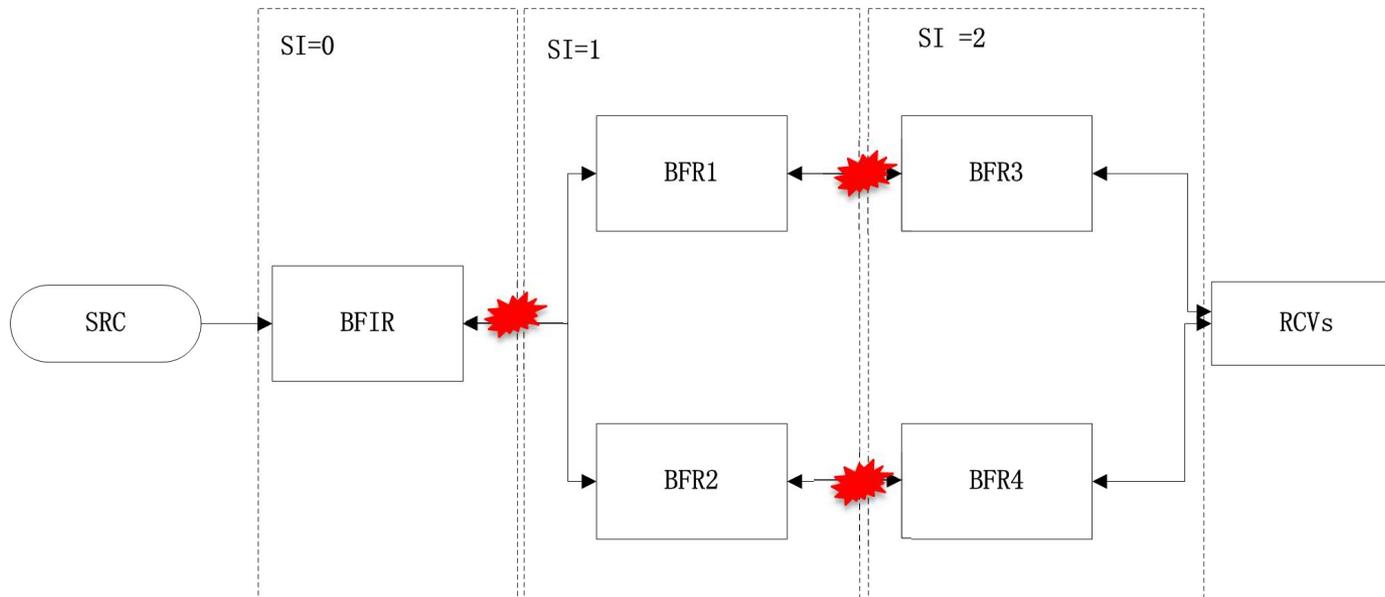
Overview



- draft-zcjh-bier-te-forwarding-00
- 00 version provided in September 2017.
- This document proposes a set of extensions to realize the BIER-TE forwarding including the assignment of BitPositions to adjacencies and the configuration of Bit Index Forwarding Table (BIFT).

Problem with the Multiple SI

- The problem is that if the adjacencies from BFIR to BFER is assigned for multiple SI and a transit BFR doesn't have the given SI for a BIER-TE packet (the BitString in BIER encapsulation is related to SubDomain, BSL and SI combination), the BIER-TE packet may be dropped in that BFR.





Motivation

- For a multicast forwarding, the controller host needs to assign lots of BitPositions and use multiple SI and BSL within the same sub-domain. The distinct SD, BSL and SI combinations MUST be mapped to more than one BitStrings and carried in different packets. Different with BIER, the BIFT related to the BitPositions which associated with a particular SD, BSL and SI combination need to be built throughout the whole network in BIER-TE.
- The BFRs need to forward the packets based on the BitString and BIFT with a SD, BSL and SI combination. The BitPositions of these adjacencies passing through BFIR to each BFER must be assigned in the same SD, BSL and SI combination to ensure the multicast flow be forwarded to the BFER within the same packet. The assignment of BitPositions and the configuration of BIFT should be taken to considerations in detail.

The process for BIER-TE forwarding



Step1: The controller host which representing the control plane of BIER-TE discovers the network topology information.

Step2: THE controller host tracks the multicast flow overlay to determine which multicast flow needs to be sent by a BFIR to which BFERs.

Step3: The controller host calculates the explicit paths from BFIR to every BFER across the BIER-TE domain.

Step4: The controller host assigns the BitPositions including SD,BSL and SI combination to the adjacencies according to the assignment method and policy. The BIFT is populated by the controller host and incremental configured Before the service deployment.

Step5: The controller host calculates the BitStrings According to the explicit paths and its related explicit SD, BSL and SI combination and pushes them into the BFIR.

Step6: Once the BIFTs and BitStrings was programmed into the data plane of BFRs by the BIER-TE controller host, they can be used to forward packets.

The Assignment of BitPositions



- This document proposes a method for the assignment of BitPosition to adjacencies and defines two types of the assignment policy of BitPositions as following.
- If the multicast flow needs to be sent from a BFIR to M BFERs along M explicit paths, the controller host MAY assign BitPositions for all adjacencies of M explicit paths with the K sets of SD, BSL and SI combinations which K M according to the assignment policy.
- EXCLUSIVE-TYPE: Each multicast flow MAY use one or more SD, BSL and SI combination exclusively.
- SHARING-TYPE: More than one multicast flows MAY share the same SD,BSL and SI combination. If the adjacencies of a path have been assigned to the same SI except some adjacencies which have not been assigned ever, the controller host SHOULD assign BP for these no- assigned adjacencies the same SI with the others. The premise is the index of the SI is enough for the assignment.

The extension of configuration for BIFT



Table 1 Extension of BIFT

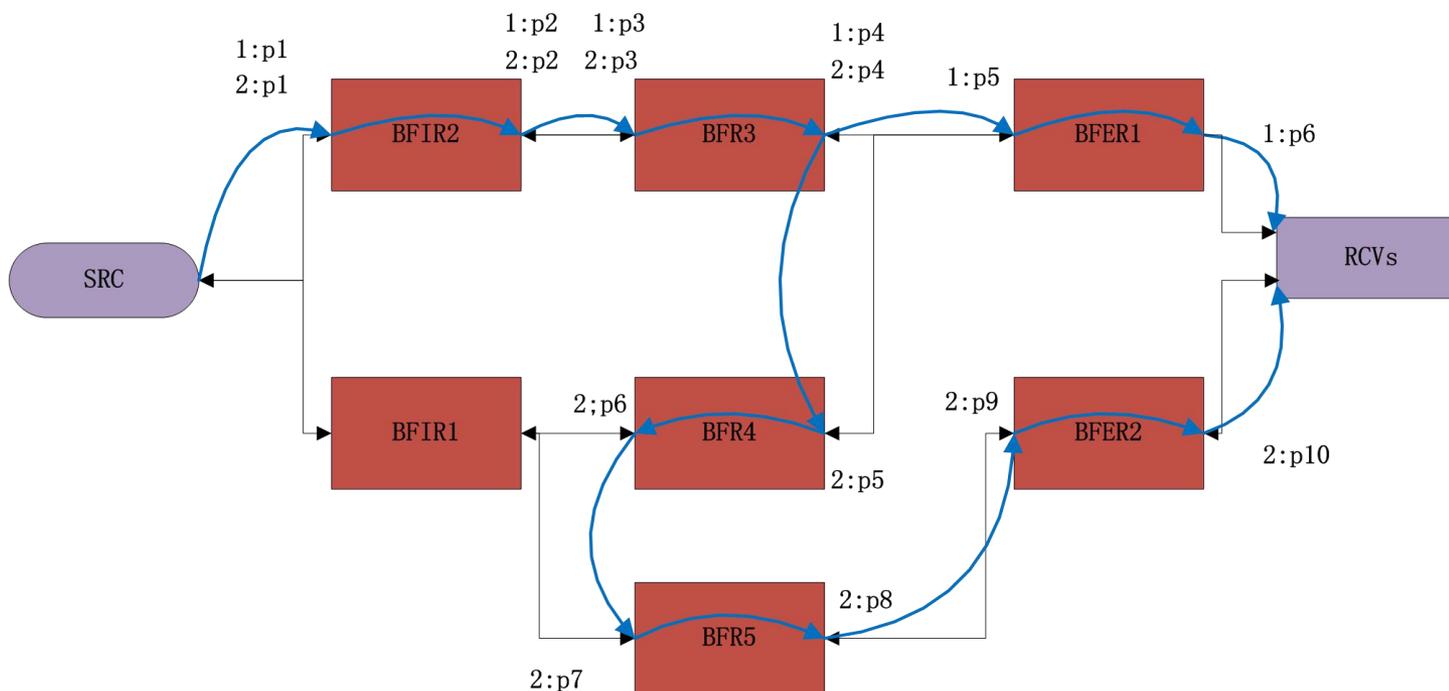
Index: BIFT-id (<SD:BSL:SI>) : BitPosition BIFT-type: BIER-TE
BIFT-id:1
BIFT-id:2
BIFT-id:3
BIFT-id:4
BIFT-id:5
BIFT-id:6
...
ID:BitStringLength

- The BIFT in one sub-domain of a BFR is a table indexed by BIFT-id:BitPosition.
- BIFT-id: indicates the ID of BIFT and MUST be assigned by BIER-TE controller host and unique throughout the BIER-TE domain and represents a BIFT.
- BIFT-type :indicates the type of BIFT including BIER and BIER-TE.
- The configuration of BIFT is not completed before the service deployment but incremental configuration based on the requirement of multicast flows.

Example1: BIRE-TE forwarding



- Multicast flow 1
- BFIR2->BFER1,BFER2
- Assign BIFT-id:BitPosition to Adjacencies for Flow 1

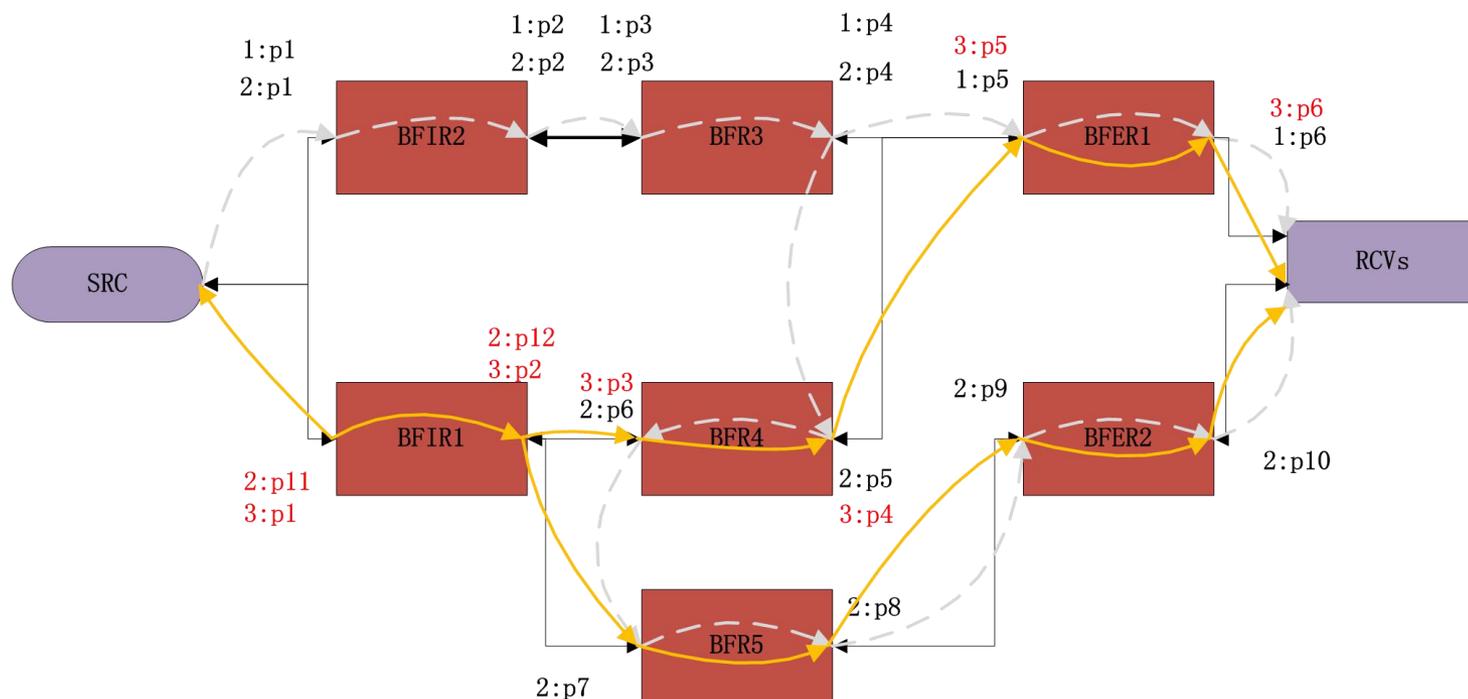


<SD=0, BSL=8, SI=0>,BIFT-id=1
 <SD=0, BSL=16, SI=1>,BIFT-id=2
 BIFT-id:BitString is
 1:0011010
 2:0000001101010100

Example2:BIRE-TE forwarding



- Multicast flow 2
- BFIR1->BFER1,BFER2
- Assign BIFT-id:BitPosition to Adjacencies for Flow 2



<SD=0, BSL=16, SI=1>,BIFT-id=2

<SD=0, BSL=8, SI=2>,BIFT-id=3

BIFT-id:BitString is

2:0000001101000000

3:00110100

Next Steps



- Discussion with related drafts' Authors
- Comments and discussion
- Request WG Adoption