

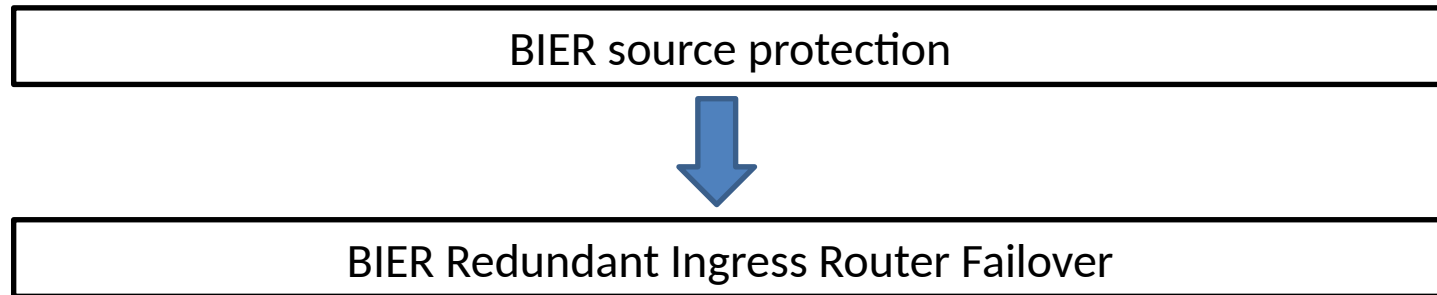
BIER Redundant Ingress Router Failover

draft-zhang-bier-source-protection-02

BIER WG
IETF110

Sandy Zhang
Gregory Mirsky
Quan Xiong
Yisong Liu


Update of 02 version



- Add reference to draft-szcl-mboned-redundant-ingress-failover, which discusses the multicast redundant ingress router failover issue.
- This draft analyses the issue in BIER domain. In the new version, more detail description is added for the issue analysis, and the node and working path failure detection.

Background

- Multicast source connects two ingress router (BFIRs) to avoid single node failure.
- BFIRs advertise the source information to all the BFERs.
- BFER learns that the multicast flow can be received from BFIR1 and BFIR2. That is, BFIR1 and BFIR2 are both UMH candidates.
- For a specific multicast flow, BFIR acts either as S-BFIR or B-BFIR. S-BFIR forwards the data flow to BFERs in response to BFERs request. B-BFIR is in the standby mode for that flow.



Standby

Cold Standby


- BFER selects a BFIR as UMH and signals to the selected BFIR to get the multicast flow.
- When BFER finds that the BFIR is down, BFER signals to the another BFIR.

Warm Standby

- BFER signals both BFIRs,
- S-BFIR forwards the flow to BFER
- B-BFIR must not forward the flow to BFER unless S-BFIR is down.

Hot Standby

- BFER signals both BFIRs
- Both BFIRs are sending the flow to the BFERs.
- BFERs must discard the duplicate flow.
- In this situation, BFIR does not distinguish the S-BFIR or B-BFIR role.



Standby functions comparison

	Cold Standby	Warm Standby	Hot Standby
BFIR	<ul style="list-style-type: none"> Regular operation forwarding flow according to the request from BFER. 	<ul style="list-style-type: none"> Takes the role of S-BFIR or B-BFIR B-BFIR MUST NOT forward flow to BFER. When S-BFIR fails or the path between S-BFIR and BFER fails, B-BFIR MUST start forwarding the flow to BFER. 	<ul style="list-style-type: none"> Need not to know the roles of S-BFIR and B-BFIR, just forwarding flow according to the request from BFER.
BFER	<ul style="list-style-type: none"> Must select a BFIR as S-BFIR to signal the flow request. Signal to the B-BFIR to request the flow when node failure or path failure between BFIR and BFER is detected. 	<ul style="list-style-type: none"> Does not select the S-BFIR or B-BFIR, just signal to both of them. 	<ul style="list-style-type: none"> Signal to both of S-BFIR and B-BFIR. Must discard the duplicate flow from B-BFIR per BFIR-ID. When S-BFIR fails or the path between S-BFIR and BFER fails, must forward the flow from B-BFIR.
BFR	<ul style="list-style-type: none"> Knows nothing about S-BFIR or B-BFIR. No duplicated flow is forwarded. 	<ul style="list-style-type: none"> Knows nothing about S-BFIR or B-BFIR. No duplicated flow is forwarded. 	<ul style="list-style-type: none"> Knows nothing about S-BFIR or B-BFIR. Duplicated flow is forwarded.

- The network administrator selects the most suitable way to deploy protection according to the network situation.

Switchover

Cold Standby


- BFER detects the failure of S-BFIR, signals to B-BFIR.
- B-BFIR forwards the flow to BFER.
- packet loss during the signaling.

Warm Standby

- B-BFIR detects the failure of S-BFIR.
- B-BFIR forwards the flow to BFER.
- packet loss during the BFIR switchover.

Hot Standby

- BFER detects the failure of S-BFIR.
- BFER stops discarding the flow from B-BFIR, BFER forwards the flow from B-BFIR.
- packet loss during the duplicate flow recognition swithing.



Failure Detection

- node failure
 - BIER overlay monitoring: Overlay period signaling stopping.
 - BIER layer monitoring:
 - B-BFIR monitors S-BFIR: IPv4/IPv6 Ping, LSP-Ping, BIER Ping.
 - BFER monitors S-BFIR:
 - Multiple P2P sessions monitoring as the same as B-BFIR monitoring S-BFIR.
 - P2MP monitoring
- path failure
 - BFER monitors S-BFIR
 - Multiple P2P sessions monitoring as the same as B-BFIR monitoring S-BFIR.
 - P2MP monitoring
 - S-BFIR monitors BFERs
 - Multiple P2P sessions monitoring as the same as B-BFIR monitoring S-BFIR.
 - P2MP monitoring

P2MP Monitoring




BIER Ping

BIER BFD


BIER ping

- BFER sends periodical ping packet to the selected UMH.
- If BFER cannot receive reply from the UMH for a period of time, BFER will treat the UMH as a failed UMH and select a new UMH.
- Not only the node failure can be detected, the path from the selected UMH to BFER can also be monitored.
 - BIER ping is defined in [\[draft-ietf-bier-ping\]](#).
 - If the path from BFER to the selected UMH is different from the path from the UMH to the BFER, the ping result may be incorrect (false negative or false positive) and the unnecessary switchover may be triggered.



BIER BFD

- The selected UMH (BFIR) sends periodical P2MP BFD control packets to all the BFERs which select the BFIR as UMH.
- BFER uses the P2MP BFD packets to monitor BFIR. If BFER cannot receive the packet for a period of time, BFER selects a new BFIR as the UMH.
- Not only the node failure can be detected, the path from the selected UMH to BFER can also be monitored.
 - BIER ping packet defined in [draft-ietf-bier-ping] is used to bootstrap the P2MP BFD session.
 - IGP and overlay extension can also be used to bootstrap the P2MP BFD session.



- Comments are welcomed ☐
- WG adoption?

Thanks!