Multicast Redundant Ingress Router Failover

draft-szcl-mboned-redundant-ingress-failover-00

MBONED WG

IETF109

Greg Shepherd
Sandy Zhang
Yisong Liu (presenter)
Ying Cheng

Multicast Domain & IR/ER

- **Multicast Domain**: forward multicast flow according to specific multicast technologies, such as PIM, BIER , P2MP TE tunnel, MLDP, etc. This domain may or may not connect the multicast source and receiver directly.
- **IR/ER**: the ingress/egress router in the multicast domain, may or may not connect the source /receiver directly. The most adjacent router to the source/receiver in the multicast domain.

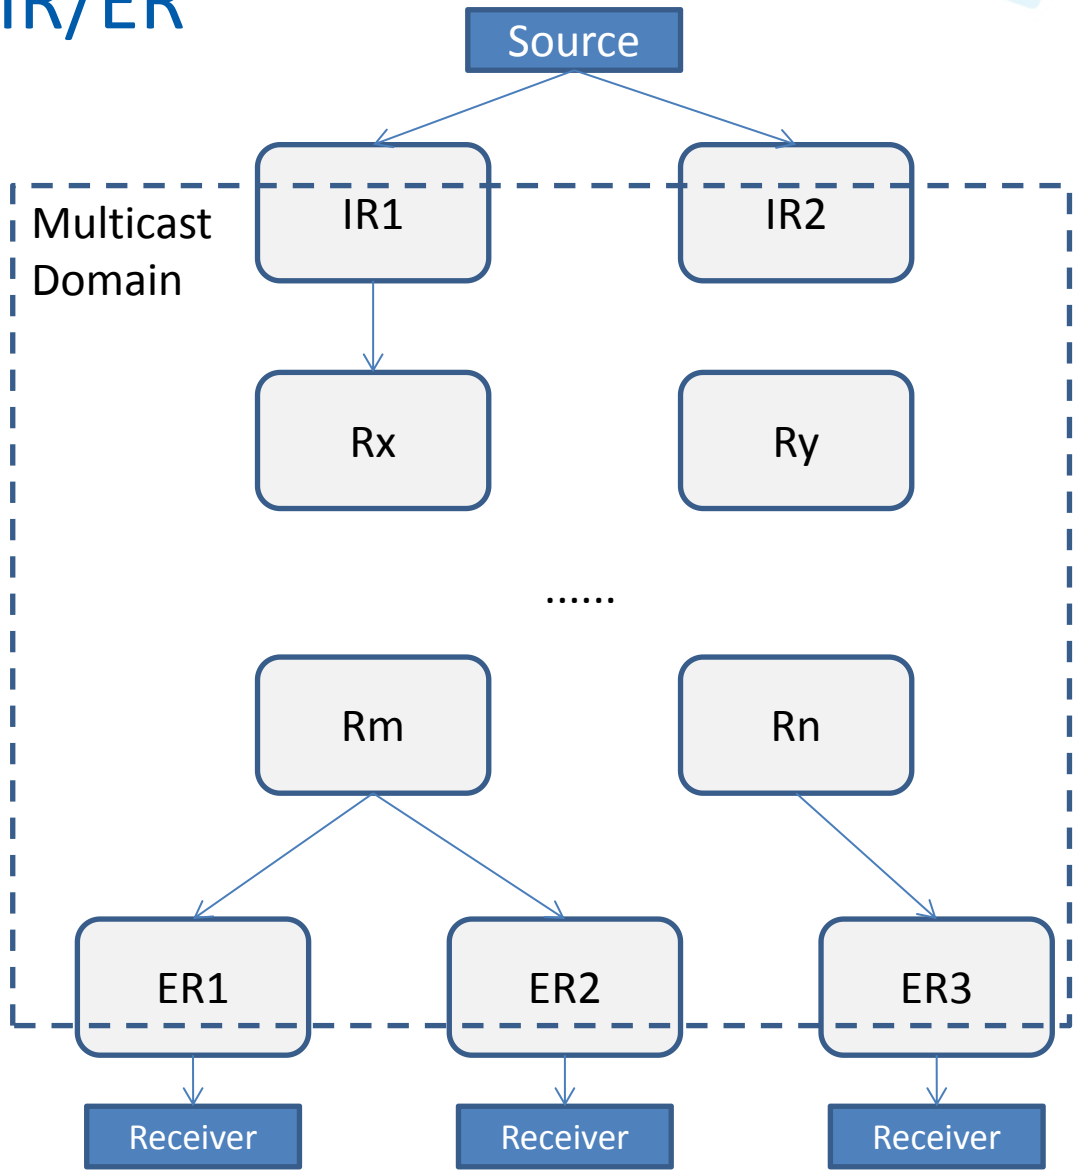

Redundant Ingress Router Failover

- A multicast source connects two IRs in order to avoid single node failure. The two IRs are all UMH candidates for ERs.
- In case the IR who is in charge of forwarding flows fails, or the path from the IR to the ERs is broken and can not be recovered, the forwarding IR switchover is needed.
- For PIM/BIER/P2MP TE tunnel/MLDP, different functions are used to built the tree or path wholly or partially.

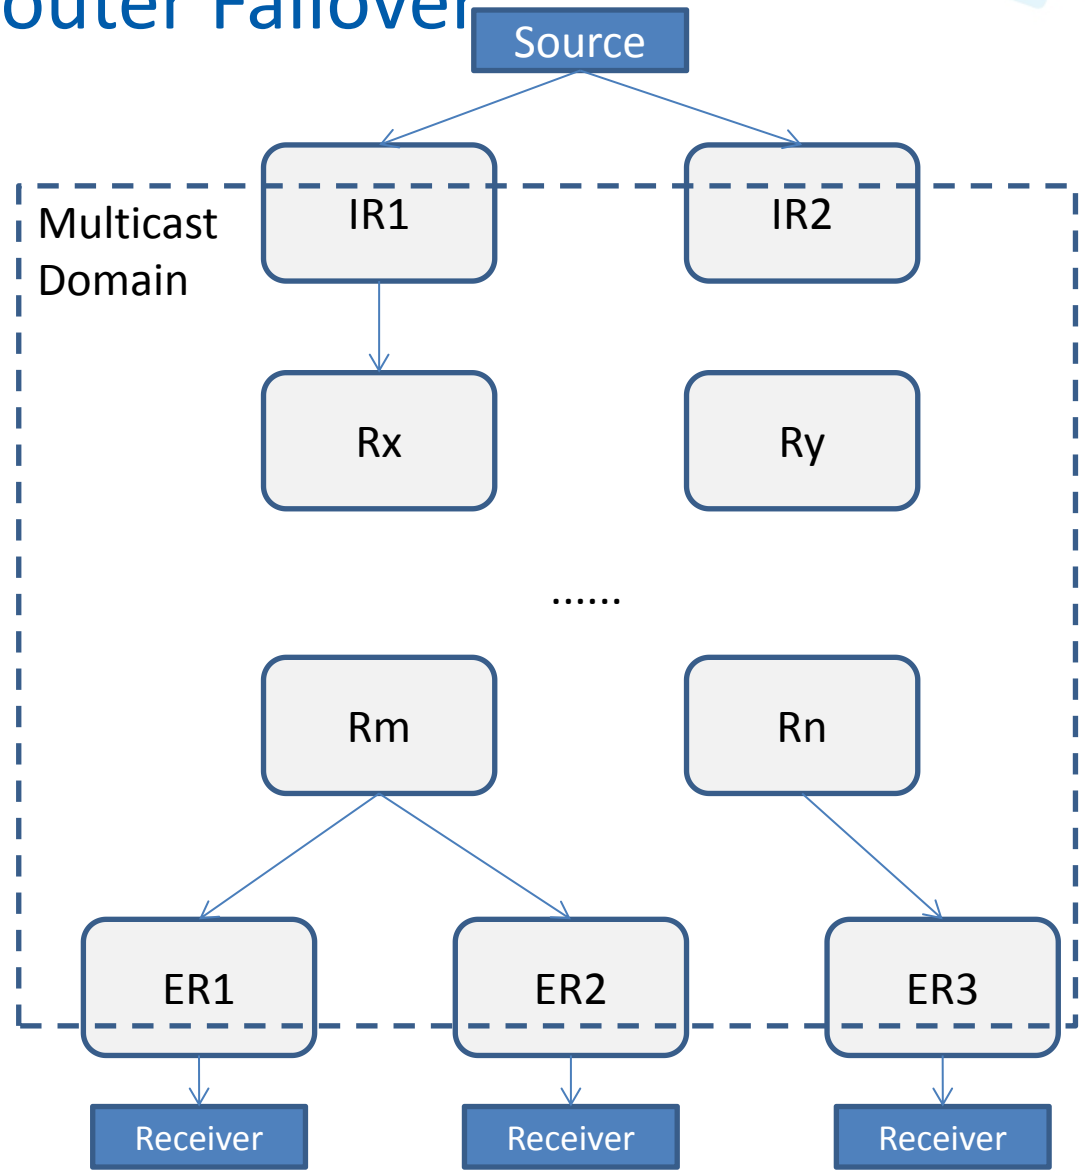

Example

IR1 is the SIR, IR1 and Rx are key nodes. When IR1 or Rx fails, there is no any other path between the IR1 and the ERs.

- PIM: Rm and Rn may choose Ry as the upstream node to send Join message to build a new tree which rooted with IR2.
- BIER: IR2 should forward to forward the flow to the ERs.
- P2MP TE tunnel/MLDP: the LSP may be rebuilt partially, or another LSP can be built in advance and replace the old one.

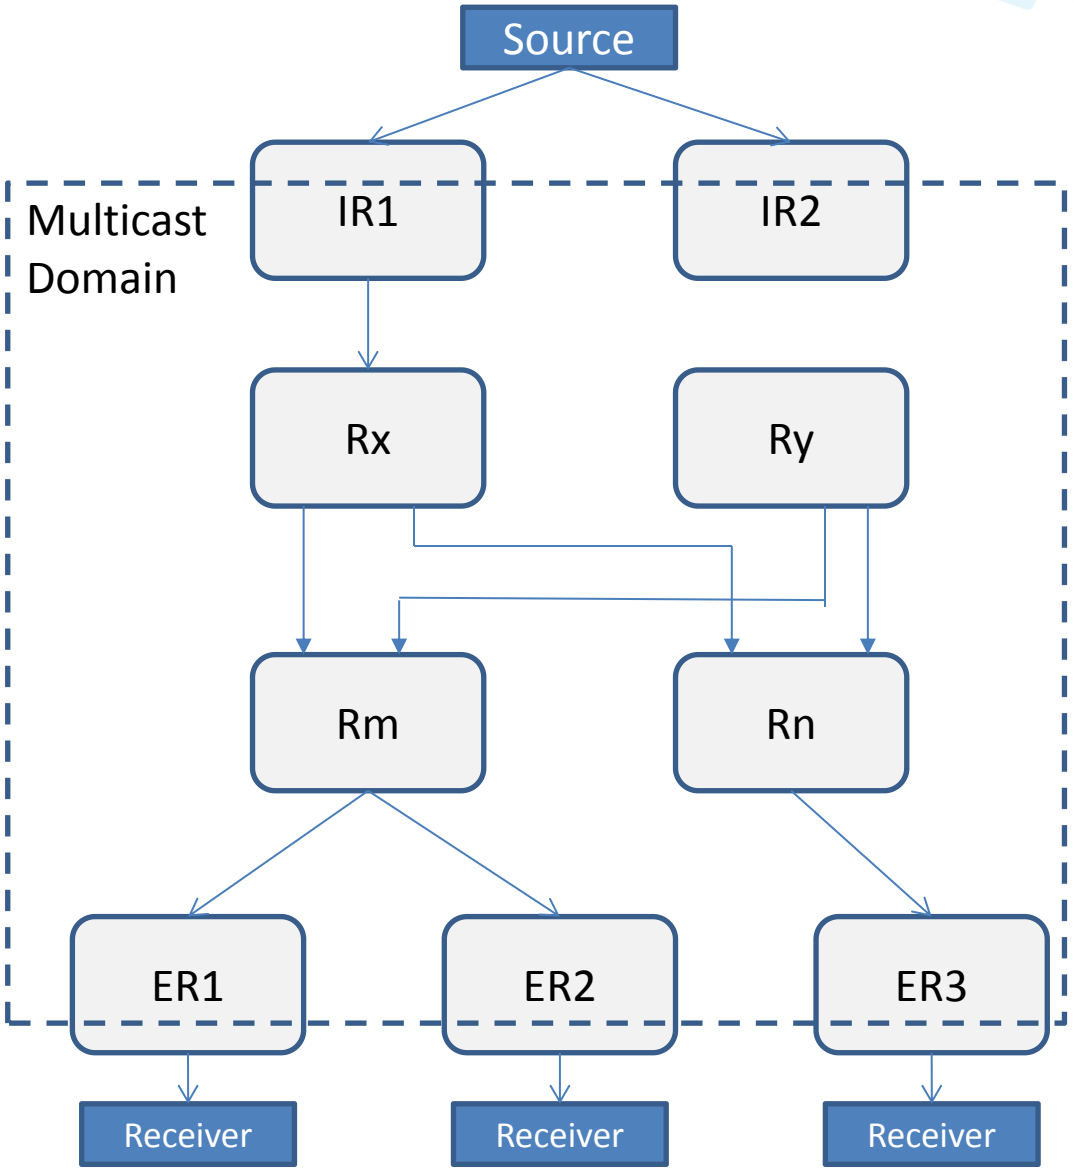

Standby Modes

Cold Standby

- ER selects an IR as Selected-IR and signals to the SIR to get the multicast flow.
- When ER finds that the SIR is down, ER signals to the BIR.

Warm Standby

- ER signals both IRs,
- SIR forwards the flow to ER
- BIR must not forward the flow to ER unless SIR is down.

Hot Standby

- ER signals both IRs
- Both IRs send the flow to the ERs.
- ERs must discard the duplicate flow.
- In this situation, IR does not distinguish the SIR or BIR role.

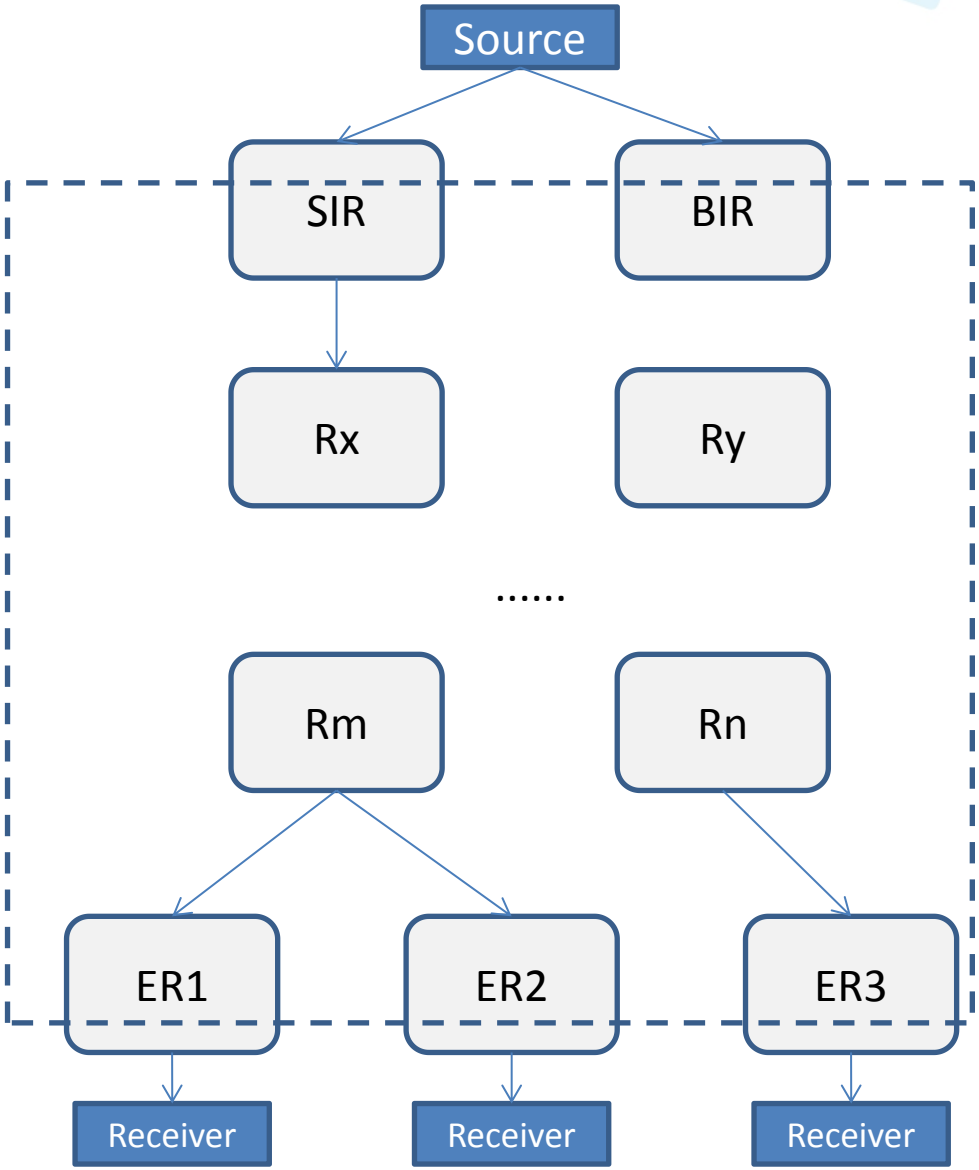

Standby functions comparison

	Cold Standby	Warm Standby	Hot Standby
IR	• forwards flow according to the request from ER.	• Takes the role of SIR or BIR • BIR MUST NOT forward flow to ER until SIR fails.	• Need not to know the roles of SIR and BIR, just forwarding flow according to the request from ER.
ER	• Must select an IR as SIR to signal the flow request. • Signal to the BIR to request the flow when SIR fails.	• Does not select the SIR or BIR, just signal to both of them.	• Signal to both of SIR and BIR. • Discard the duplicate flow from BIR until SIR fails.
R	• Knows nothing about SIR or BIR. • No duplicated flow is forwarded.	• Knows nothing about SIR or BIR. • No duplicated flow is forwarded.	• Knows nothing about SIR or BIR. • Duplicated flow is forwarded.
Conclusion	• The easiest way to implementated, but it takes the longest converge time.	• Takes the middle packet loss and converge time, but it's hard for BIR to know the failure between SIR and ERs.	• Takes the most less packet loss, but there is duplicated packet forwarding in the domain, more bandwidth is occupied.

- The SIR failure includes the SIR node failure, and the path failure between SIR and ERs.

- Comments are welcomed ☺

Thanks!