Bit Indexed Explicit Replication – A Stateless Multicast Architecture

Nagendra Kumar Nainar

NANOG72

Agenda

- Multicast Architecture Challenges
- Introduction to BIER
- BIER Control plane behavior
- BIER Data plane Encapsulation
- Packet forwarding semantic
- BIER Use cases
- Standardization Efforts

Multicast Architecture - Challenges

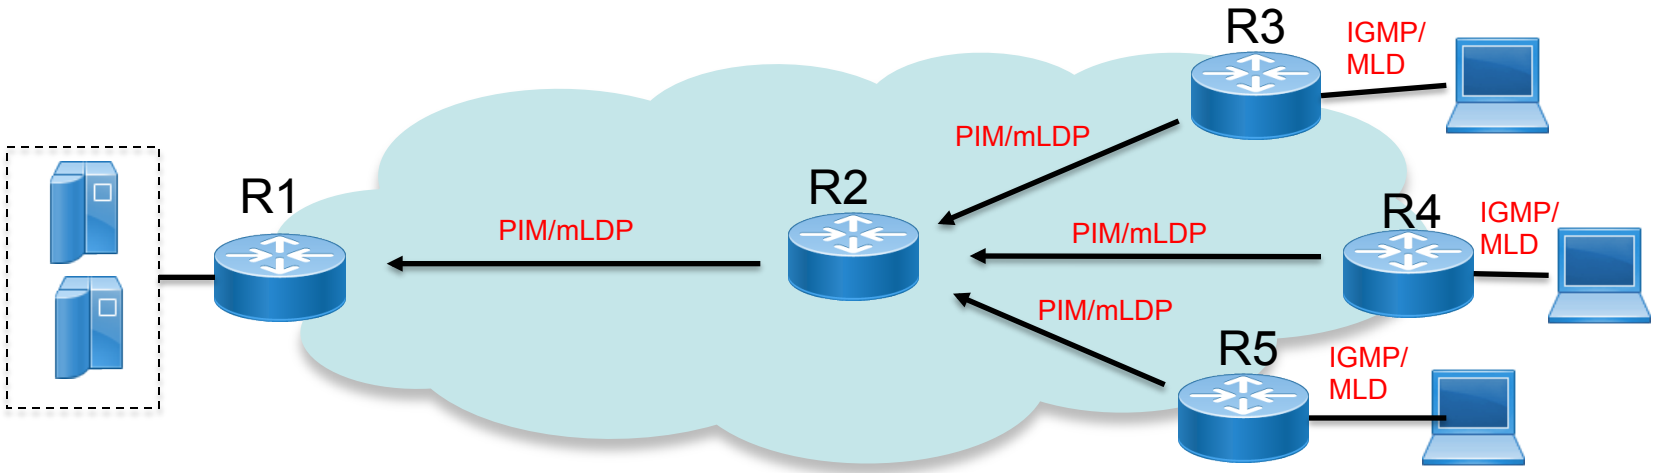

- Scalability
  - Control plane and Data plane state entries created on all transit nodes.
- Inefficient Load sharing
  - No Path control
- Poor Path Resiliency

BIER Epiphany

- We call this architecture as Bit Indexed Explicit Replication (BIER).
- **Goal**
  - Carry state entry directly in the packet header.
  - No state entries in any transit nodes
- **Control Plane semantic**
  - Assign unique bit position for each receivers
  - Propagate the info within the domain
- **Data Plane Semantic**
  - Encode the set of receivers as bit string in packet header
  - Replicate based on BIER forwarding table

Bit Position Assignment

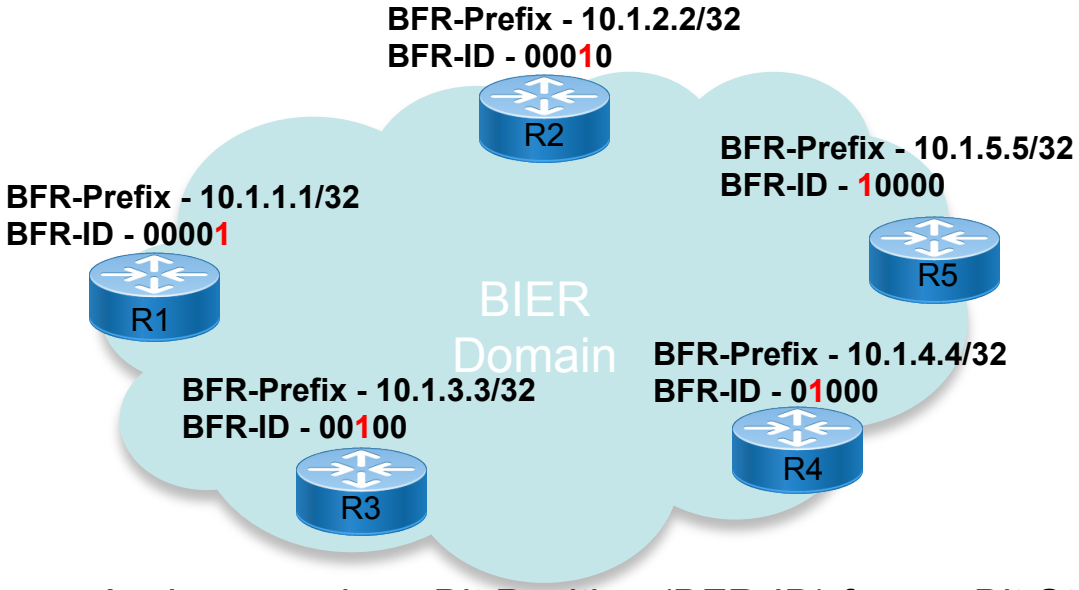

BFR – BIER Forwarding Router
BFR-Prefix – Reachable IP Address
BFR-ID – Unique Bit Position
BFIR – BIER Forwarding Ingress Router
BFER – BIER Forwarding Egress Router

BitString

- Assigns a unique Bit Position (BFR-ID) from a Bit String to each BFR within the domain.
  - String size can be 256, 512, 1024,2056 etc.
  - Bit String of size 256 can accommodate 256 BFER
- Map the BFR-ID to BFR-Prefix (locally reachable address).
  - Preferably local loopback address

Bit Position Advertisement

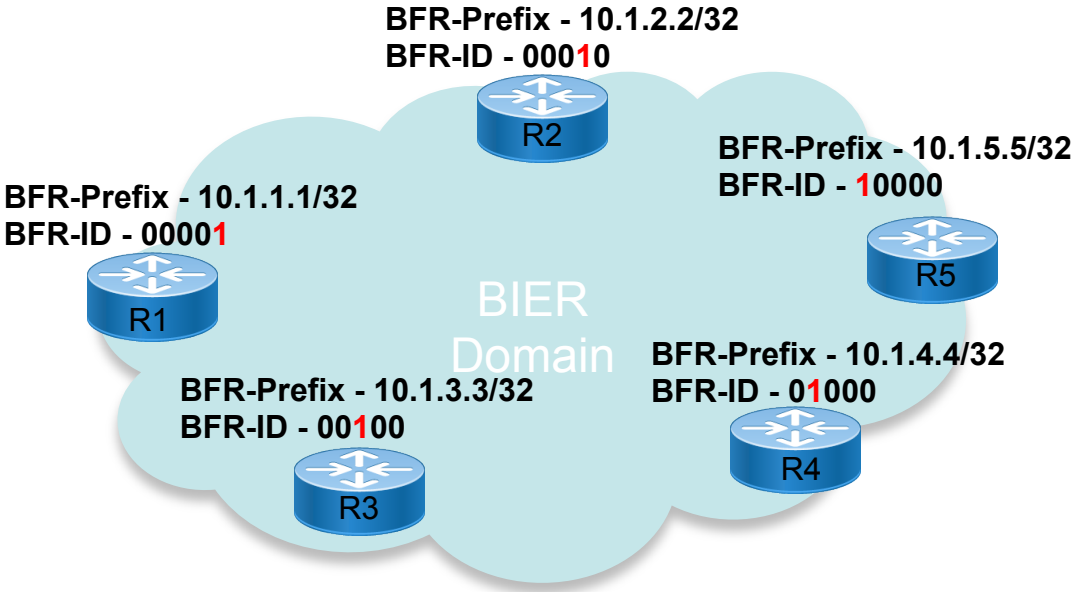

- Each BFR flood the “BFR-ID to BFR-Prefix” mapping to other nodes within the domain
  - ISIS, OSPF, BGP

Bit Index Forwarding table – Path Computation

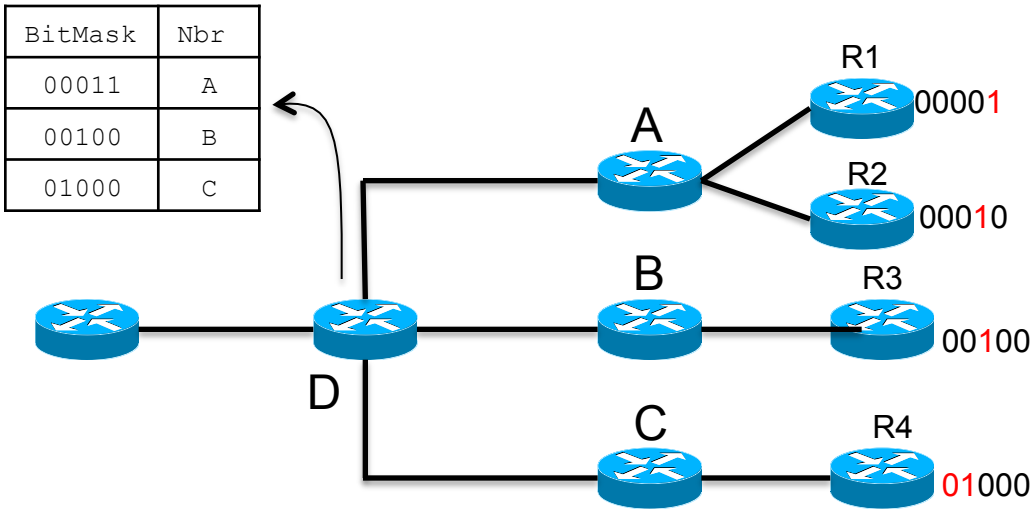

- Transit nodes compute the shortest path to each BFR-Prefix.
- Populate the forwarding table as below:
  - Identify the set of BFR-Prefix reachable to same nexthop
  - Perform “OR” operation on all the BFR-IDs reachable via same nexthop

Bit Index Forwarding Table

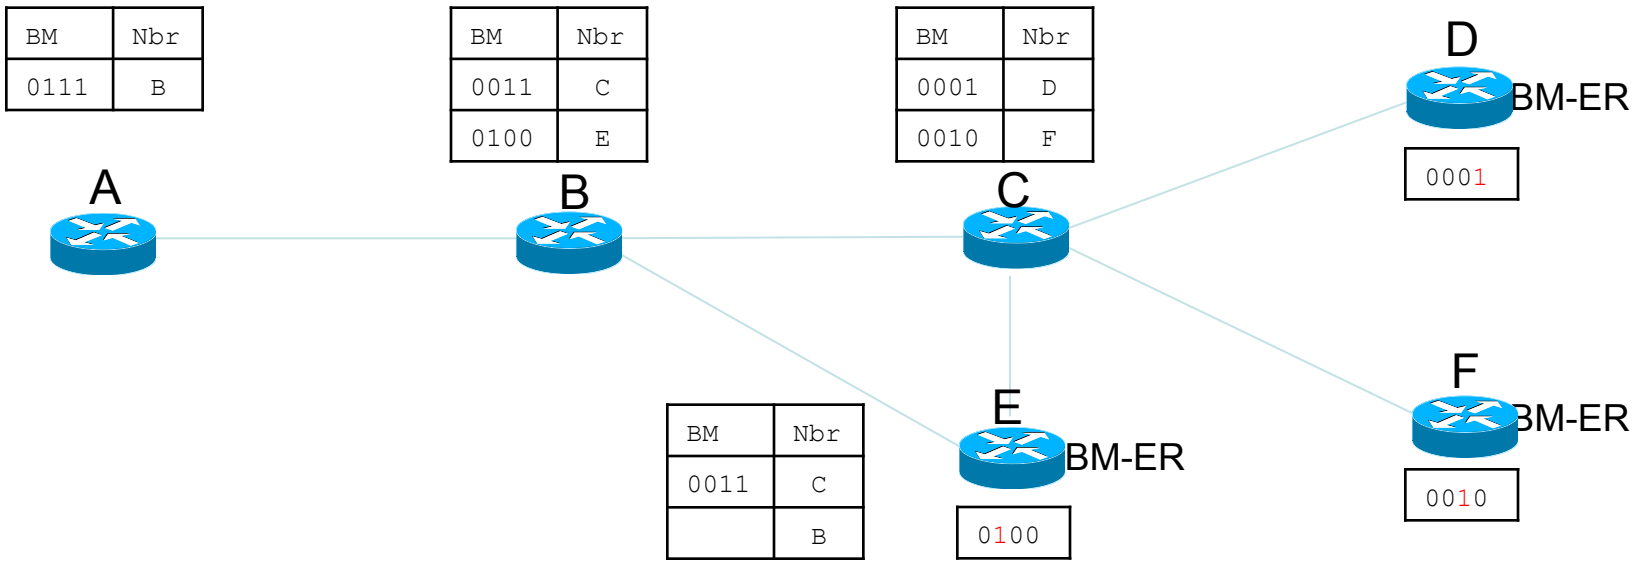

D, F and E advertise their Bit positions in the IGP (flooded).
Based on shortest path route to RID, the Bit Index Forwarding Table is created

BIER - Data Plane Encapsulation

BIER Header

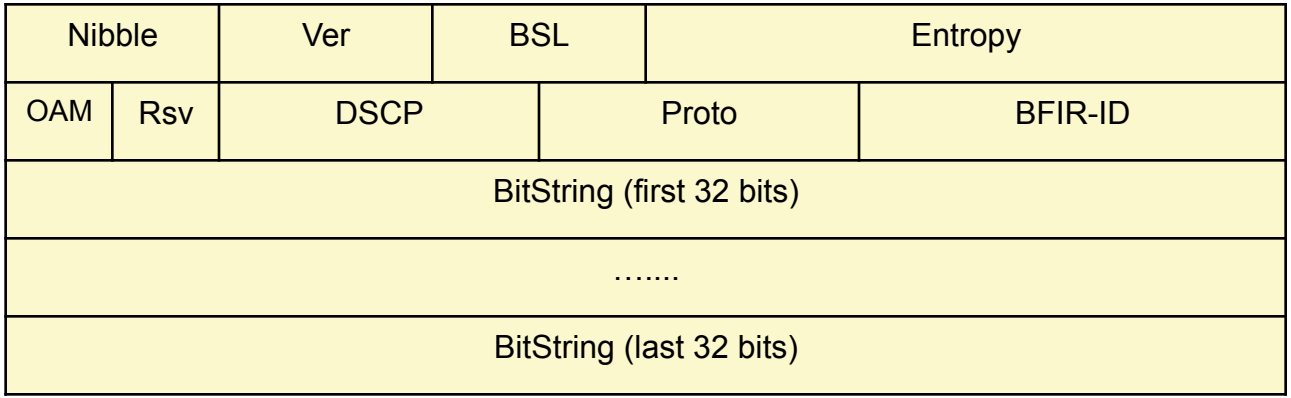

draft-ietf-bier-mpls-encapsulation

- Ether type helps differentiate the Layer PDU
- Two ways are under standardization.
  - Define a new Ether type
  - MPLS Label

MPLS encapsulation

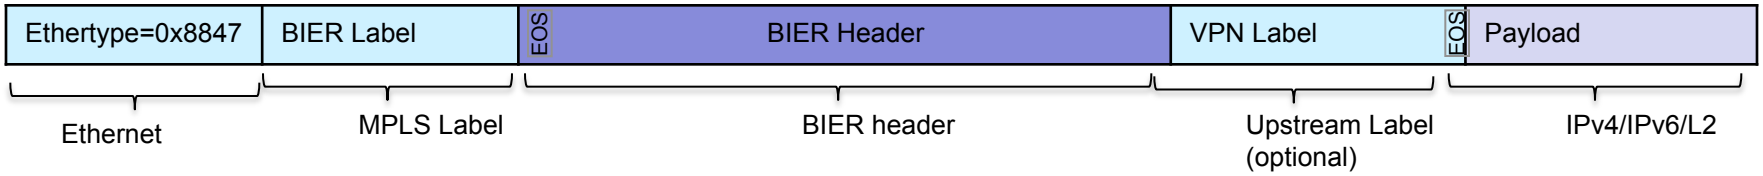

- MPLS label as a context Identifier
- MPLS label identifies the below:
  - BIER encapsulation
  - Bit String Size
  - Sub Domain ID

BIER-MPLS label Advertisement

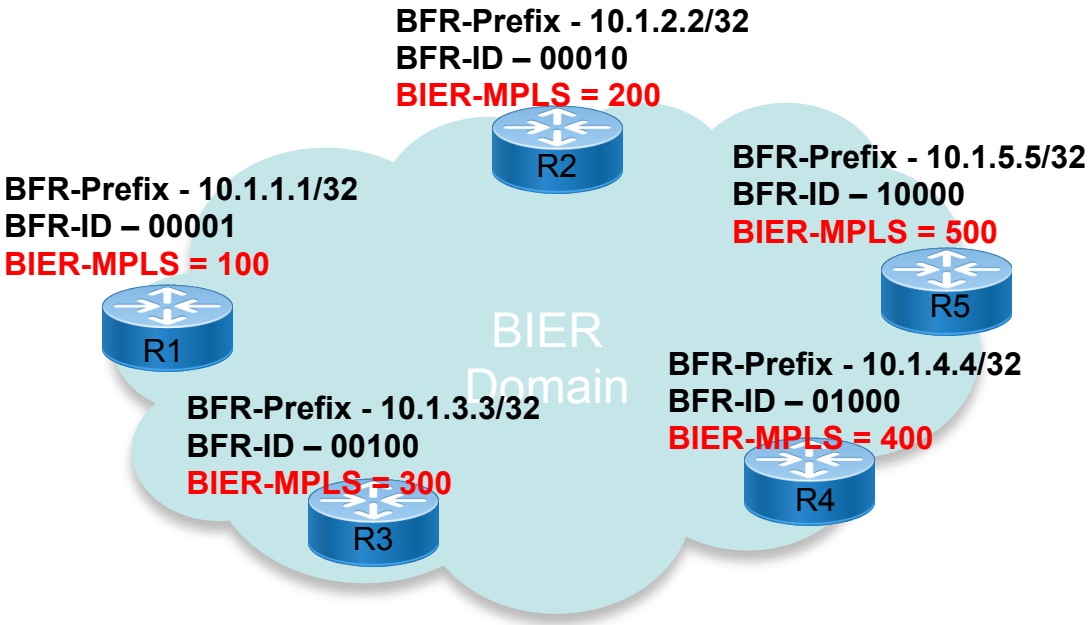

- Assigns a locally unique MPLS label and flood within the domain using IGP extensions

BIER – Packet Forwarding

BIER Forwarding Semantic

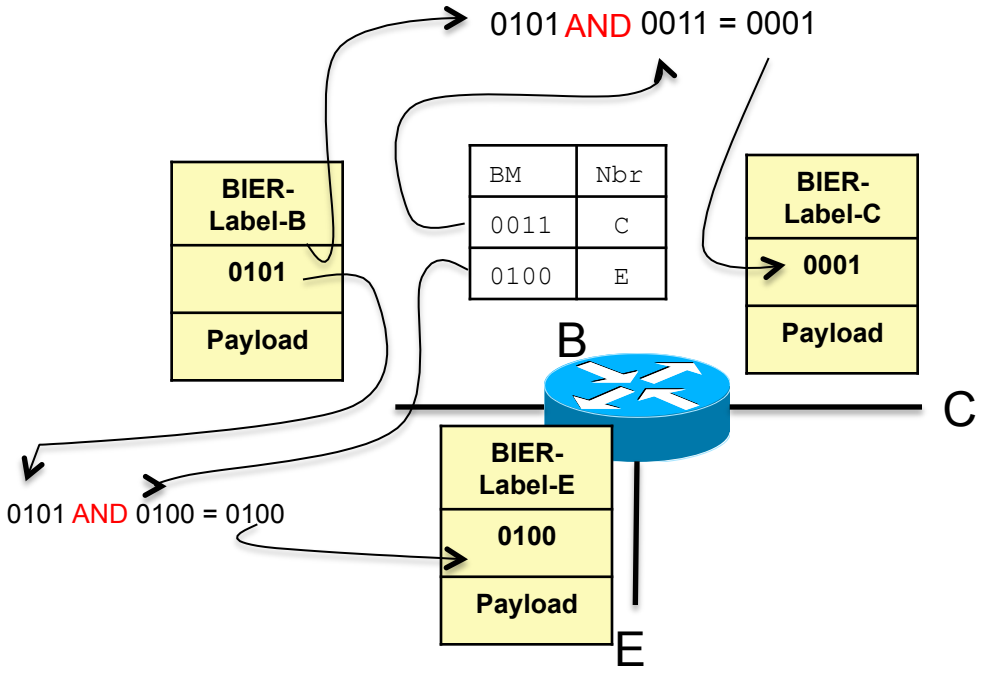

1. Determine Bit String from received packet
2. Use the LSB position set to 1.
3. Identify the first match in forwarding table
4. Perform “AND” operation between header string and table string.
5. Rewrite the header and forward to neighbor.

BIER – Packet Forwarding

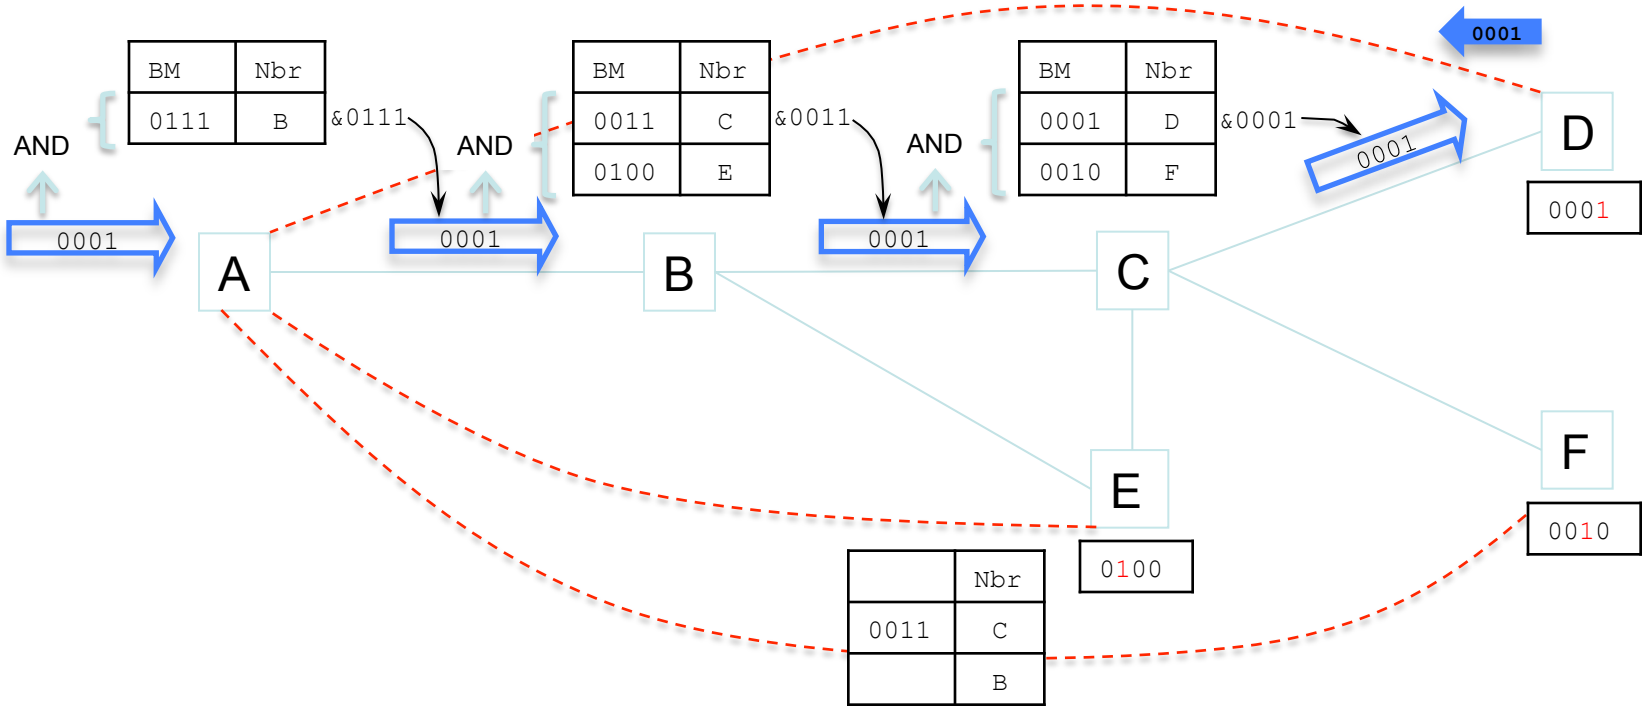

BIER – Packet Forwarding

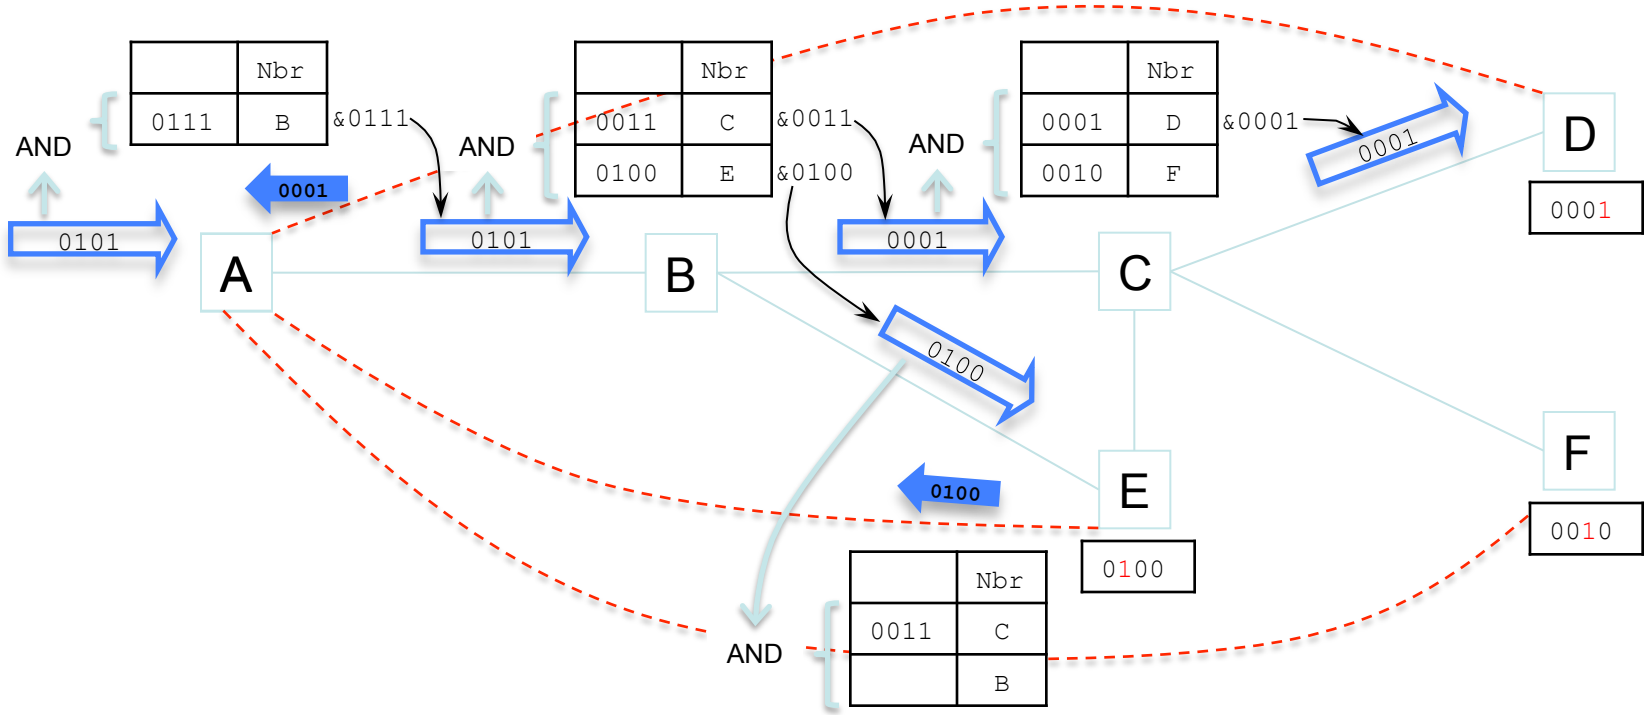

BIER Characteristics

- Stateless
  - No per flow state entries created.
  - Less Network churn.
- Duplication Avoidance
  - Bitwise AND operation avoids duplicate
- Flexible
  - Architecture can accommodate different Bit string sizes.
- Scalable

BIER Sub Domains

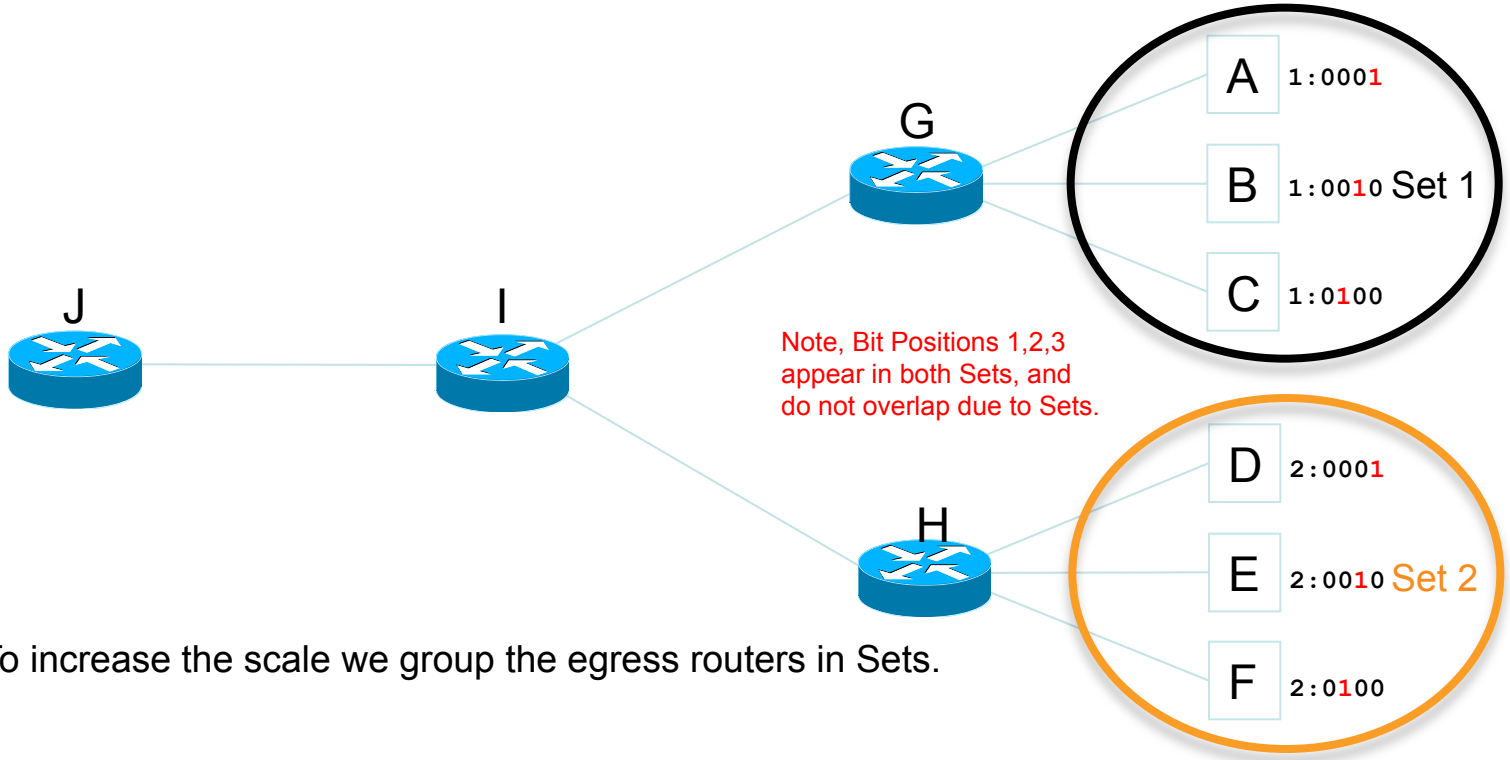

To increase the scale we group the egress routers in Sets.

BIER Set Identifier

BFR-ID = 1027

When Bit String Size is 2056, SI = 0, Bit Position = 1027
When Bit String Size is 1024, SI = 1, Bit Position = 3
When Bit String Size is 512, SI = 2, Bit Position = 3
When Bit String Size is 256, SI = 4, Bit Position = 3

- Set Identifier (SI) is identified based on the BFR-ID and Bit String Size.

BIER Sets

Set	BM	Nbr
1	0111	I
2	0111	I

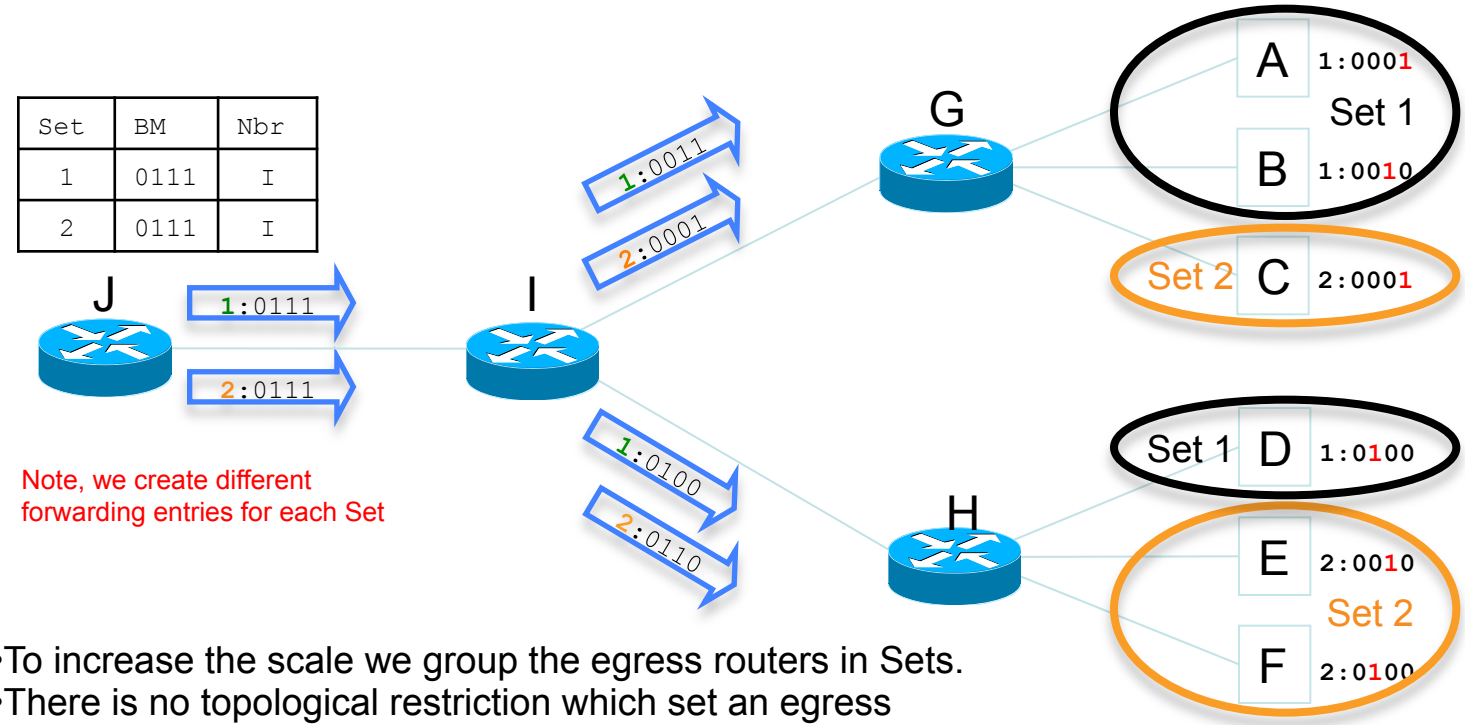

Note, we create different forwarding entries for each Set

- To increase the scale we group the egress routers in Sets.
- There is no topological restriction which set an egress belongs to

Underlay Protocol Extensions

- IGP and BGP protocols are extended to carry:
  - Sub Domain ID (Set Identifier)
  - BFR-ID
  - BFR-Prefix
  - BIER-MPLS Label
- OSPF uses Extended Prefix Opaque LSA
- ISIS uses TLVs 235, 237, 135, 236.

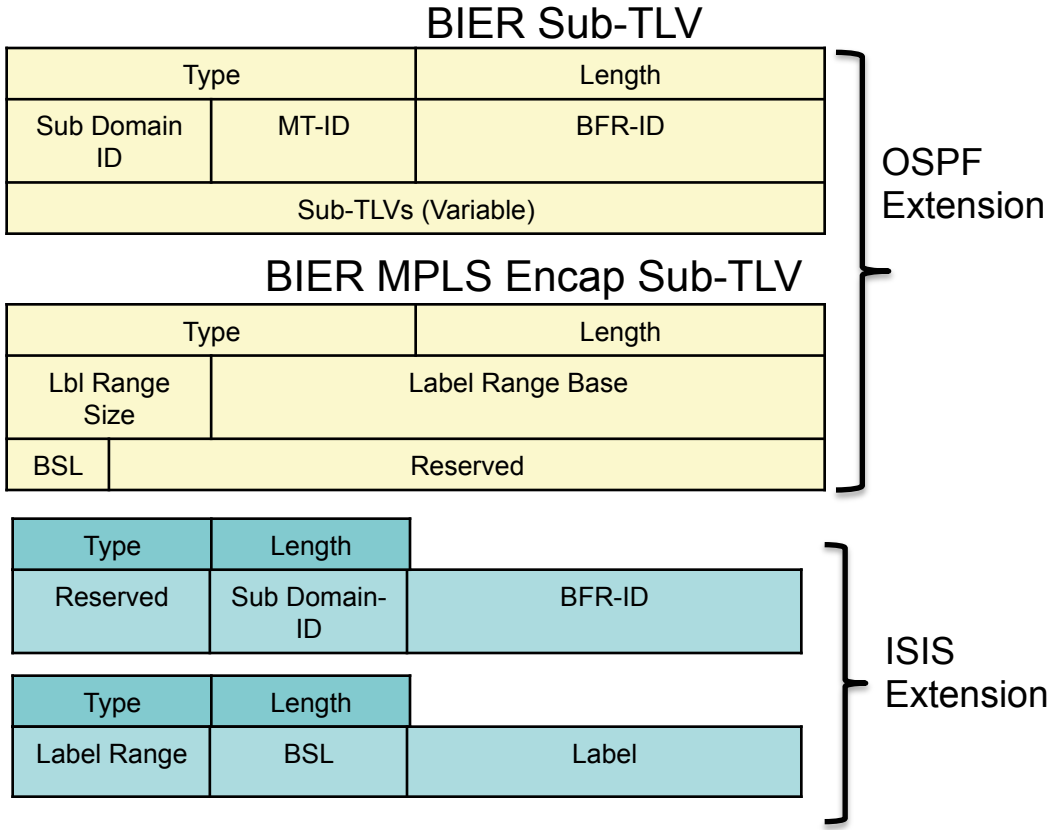

BIER Use cases

MVPN over BIER

- BIER replaces PIM, mLDP, RSVP-TE or IR in the core.
- BIER represents a full mesh (P2MP) connectivity between all the PE's in the network.
- There is no need to explicitly signal any MDT's (or PMSI's).
- With MVPN there are many profiles,
  - This is partly due to the tradeoff between 'State' and 'Flooding'.
  - Different C-multicast signaling options.
- MVPN over BIER, there is one profile.
  - BGP for C-multicast signaling.
- No need for Data-MDTs.

Traditional MVPN

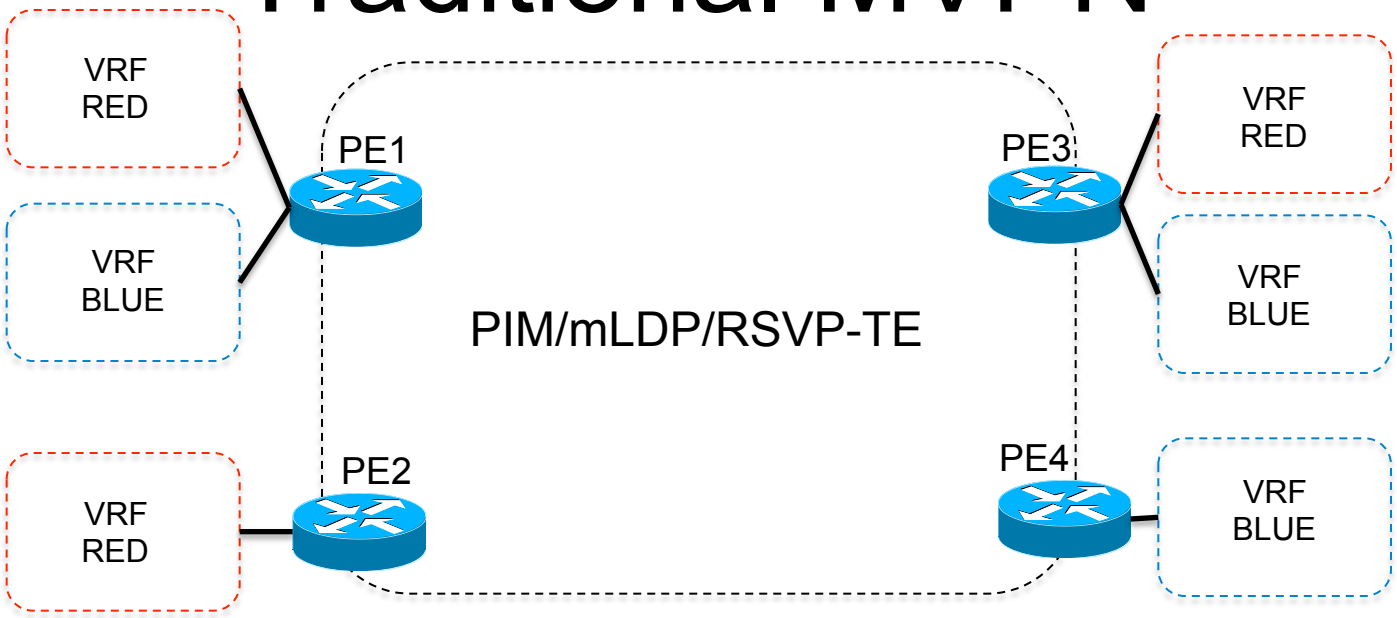

- Core to be enabled with one of the below:
  - PIM, mLDP, RSVp-TE
- Per VRF MDT tree is required to be instantiated.
- Data MDT required for optimal forwarding

MVPN over BIER

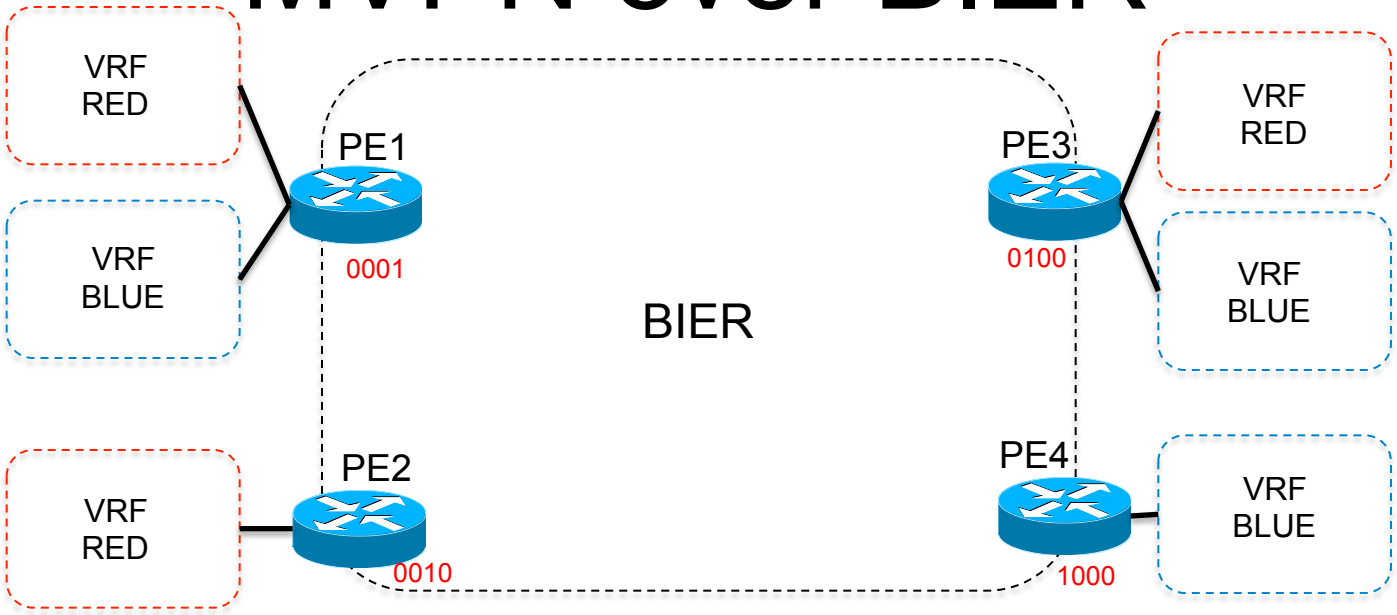

- Re-Use BGP-MVPN AFI for overlay signaling
- No per VRF MDT or tree to be instantiated
- No Data MDT required for optimality

MVPN over BIER

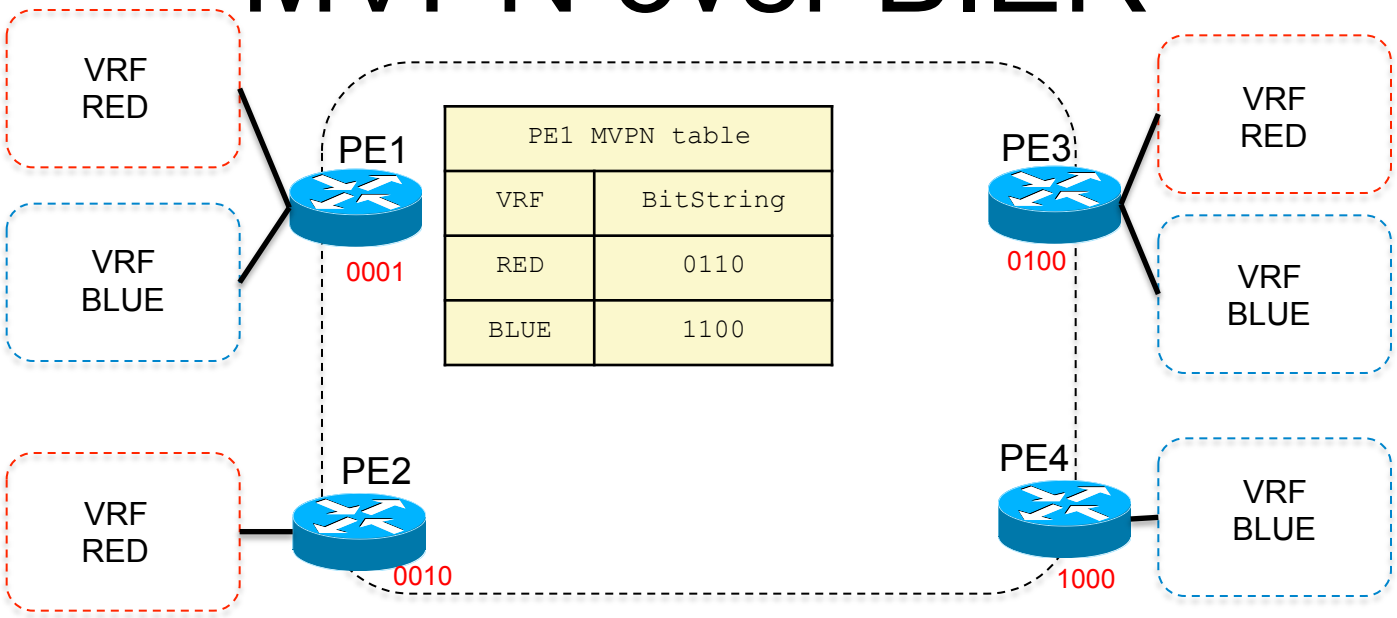

- Each PE creates the Bit String for each VRF based on overlay signaling.
- No per VRF state entries or Data MDT required

Brown field deployment

BM	Nbr	Intf
0011	C	tunnel1
…	..	..

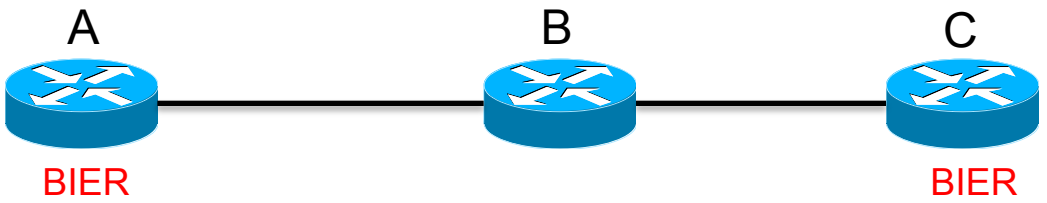

- Uses IGP extension to determine BIER capability of neighbot
- Uses unicast tunnel between BIER nodes
  - Traditional forwarding on unsupported nodes.

Brown field deployment

BM	Nbr	Intf
0011	C	tunnel1
1100	D	tunnel2

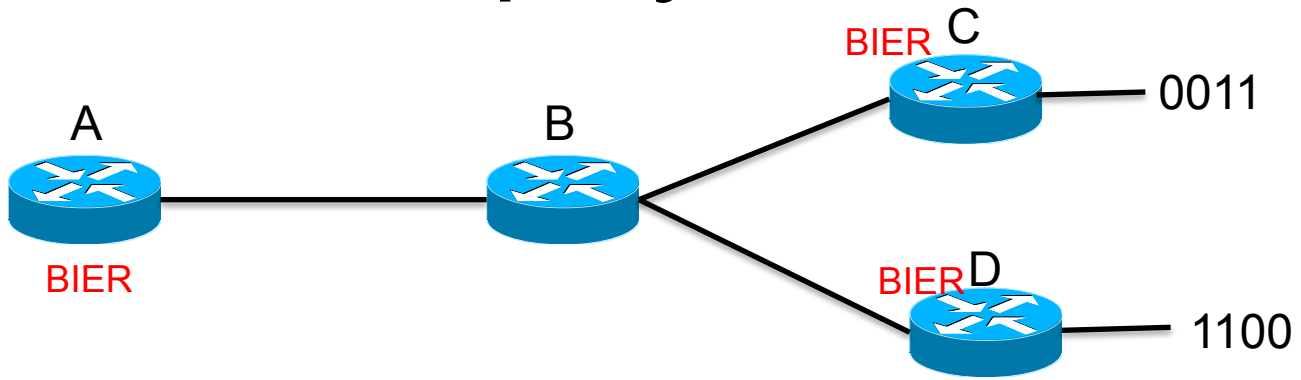

- BIER node uses different tunnel to each neighbors over traditional node.
- Rewrite the bit string and unicast to each BIER neighbor

Standardization Efforts

IETF

- BIER was proposed as BoF in IETF
  - Nov 2014 (Hawaii)
- Industry interest resulted in BIER Working Group
  - BIER (bier@ietf.org)
- BIER Architecture is published as RFC8279
- Multi Vendor Collaboration

IETF drafts

IETF Draft	Description	Status
draft-ietf-bier-architecture	BIER Architecture	Waiting for Publication
draft-ietf-bier-encapsulation-mpls	MPLS Dataplane Encapsulation	Waiting for Publication
draft-ietf-ospf-bier-extensions	OSPF Extension	Stable
draft-przygienda-bier-isis-ranges	ISIS Extension	Stable
draft-eckert-bier-te-arch	BIER Traffic Engineering Arch	Under Progress
draft-ietf-l3vpn-mvpn-bier	BIER MVPN	Under Progress
draft-ietf-bier-ping	BIER OAM	Under Progress

BIER - Advantages

- Packets forwarded via BIER follow the unicast path towards the receiver, inheriting unicast features like FRR and LFA.
- There is no per multicast flow state in the network.
- Multicast convergence is as fast as unicast, there is no multicast state to re-converge, signal, etc.
- Nice plugin for SDN, its only the ingress and egress that need to exchange Sender and Receiver information.
- The core network provides a many-2-many connectively between all BIER routers by default following the IGP.
- No Multicast control protocol in the network.

Q&A