Stateless Multicast with Bit Indexed Explicit Replication (BIER)

Alvaro Retana (aretana@cisco.com)
Distinguished Engineer, Cisco Services

Slides by: IJsbrand Wijnands

Bit Indexed Explicit Replication (BIER)

- Only encode the end-receivers in the packet header.
  - Not the intermediate nodes.
- Assign end-receivers a Bit Position from a Bit String.
  - The smallest identifier possible.
- Encode the Bit String in the packet header.
  - Using some sort of encapsulation.
- Create a Bit Forwarding Table on all BIER nodes to allow multicast packet forwarding using the Bit String in the packet.
  - Derived from the RIB, SPF based.

IETF

- The BIER idea was presented in a BOF at the IETF in Hawaii.
  - November 2014.
- A new BIER Working Group has been formed (bier@ietf.org)

draft-ietf-bier-problem-statement-00

draft-ietf-bier-architecture-00

draft-ietf-bier-encapsulation-mpls-00

draft-ietf-bier-use-cases-00

draft-ietf-l3vpn-mvpn-bier-00

draft-ietf-ospf-bier-extensions-00

draft-przygienda-bier-isis-ranges-01

draft-eckert-bier-te-arch-00

draft-xu-idr-bier-extensions-00

Solution Overview

Basic Idea BIER

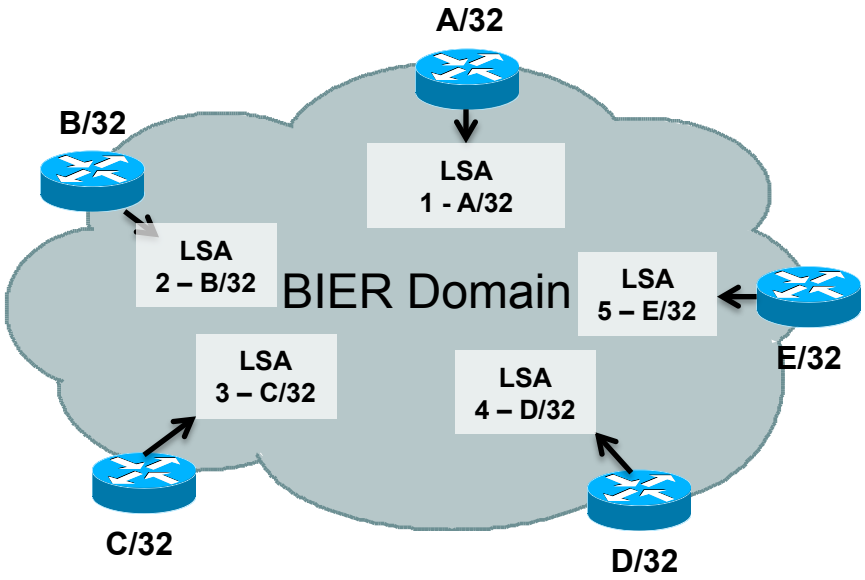

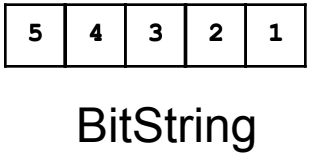

1. Assign a unique Bit Position from a BitString to each BFER in the BIER domain.
2. Each BFER floods their Bit Position to BFR-prefix mapping using the IGP (OSPF, ISIS)

Basic Idea BIER

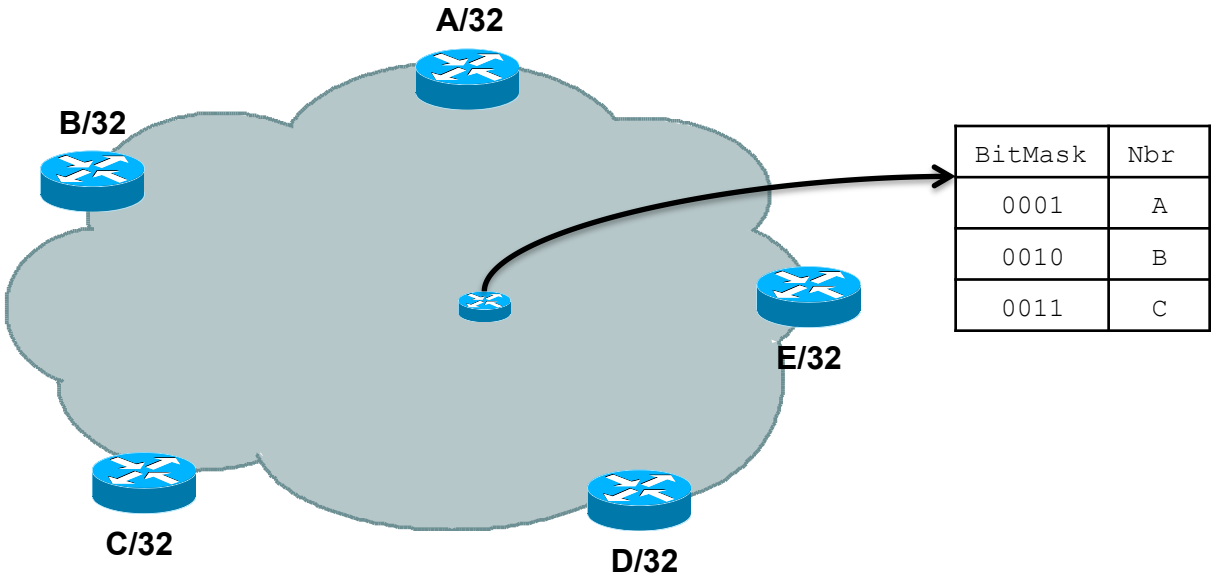

BitMask	Nbr
0001	A
0010	B
0011	C

1. Assign a unique Bit Position from a BitString to each BFER in the BIER domain.
2. Each BFER floods their Bit Position to BFR-prefix mapping using the IGP (OSPF, ISIS)

Bit Index Forwarding Table

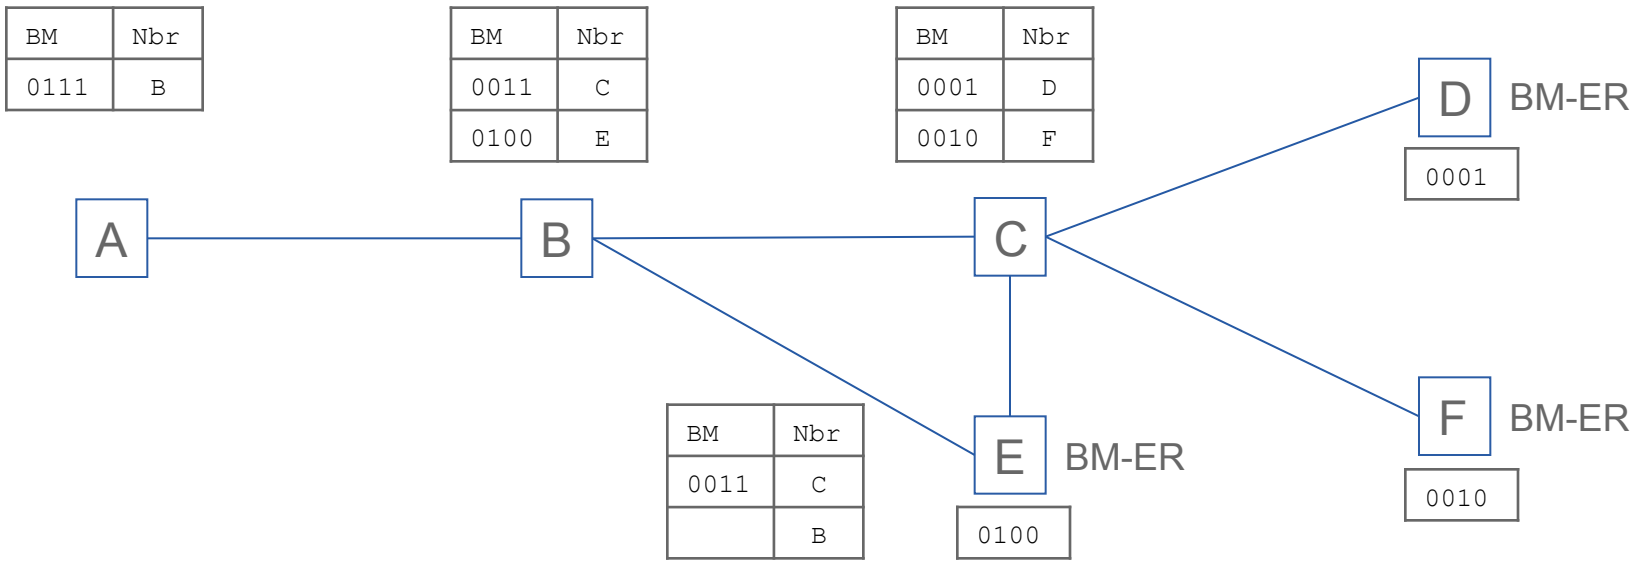

- D, F and E advertise their Bit positions in the IGP (flooded).
- Based on shortest path route to RID, the Bit Mask Forwarding Table is created

Forwarding Packets

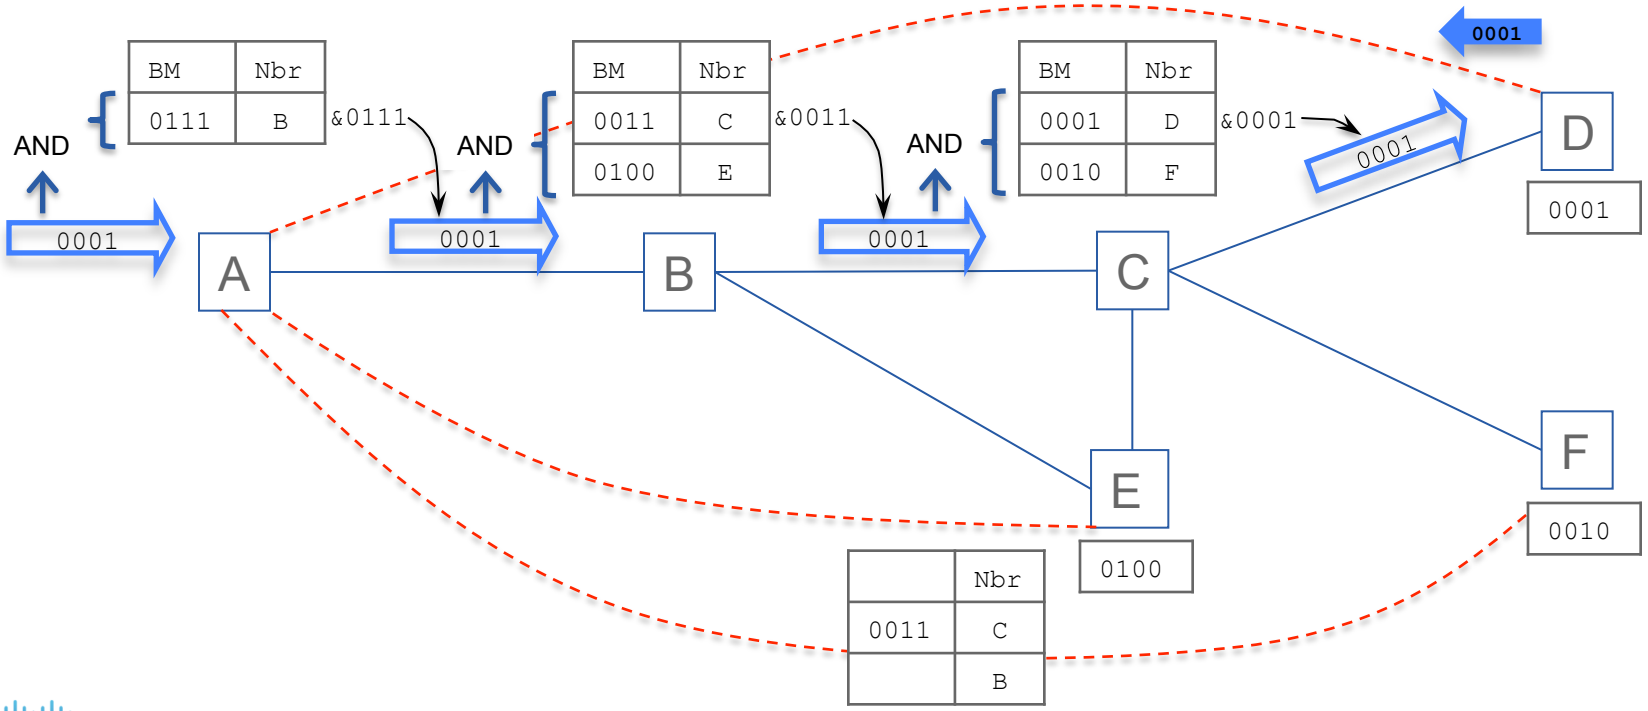

Forwarding Packets

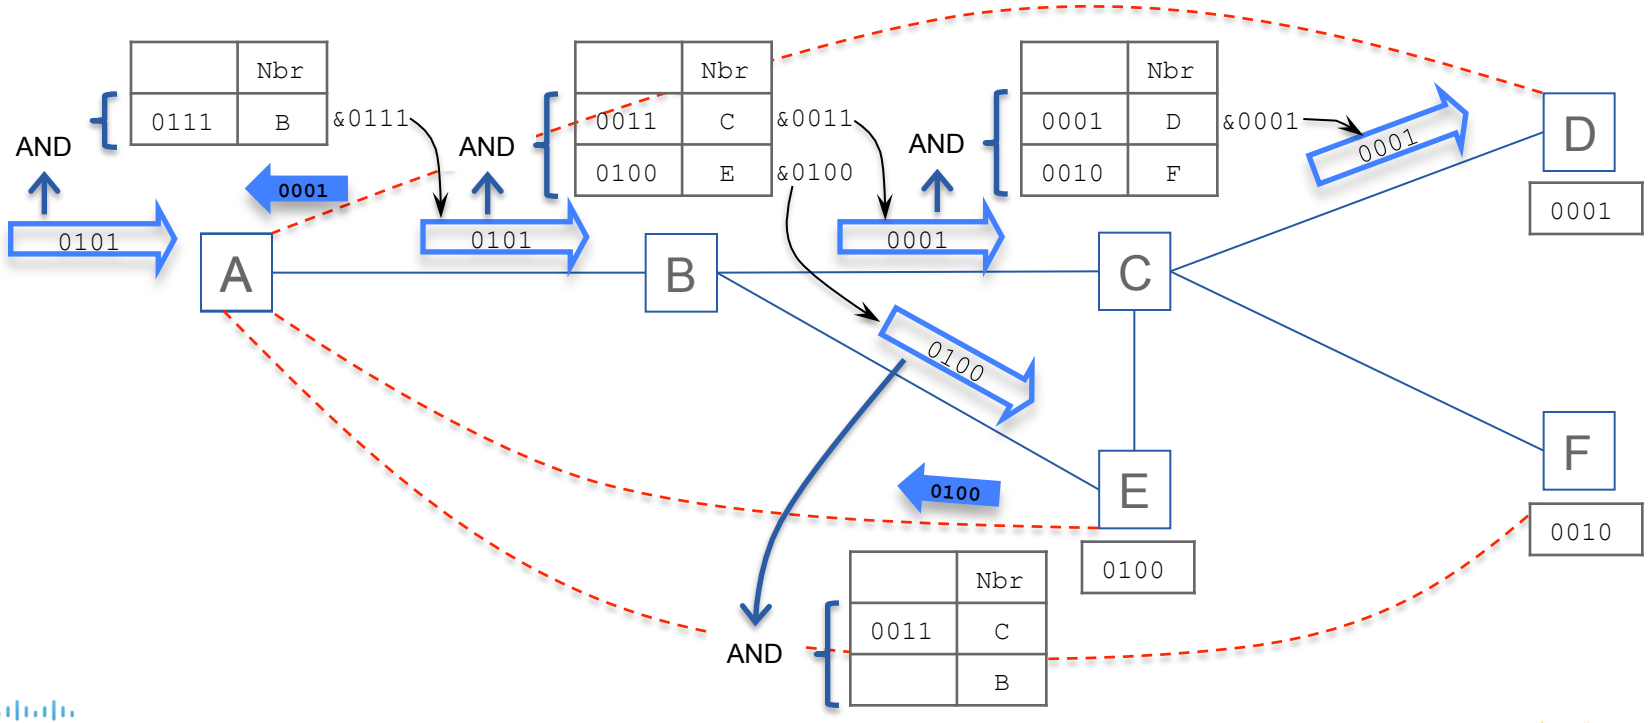

Forwarding Packets

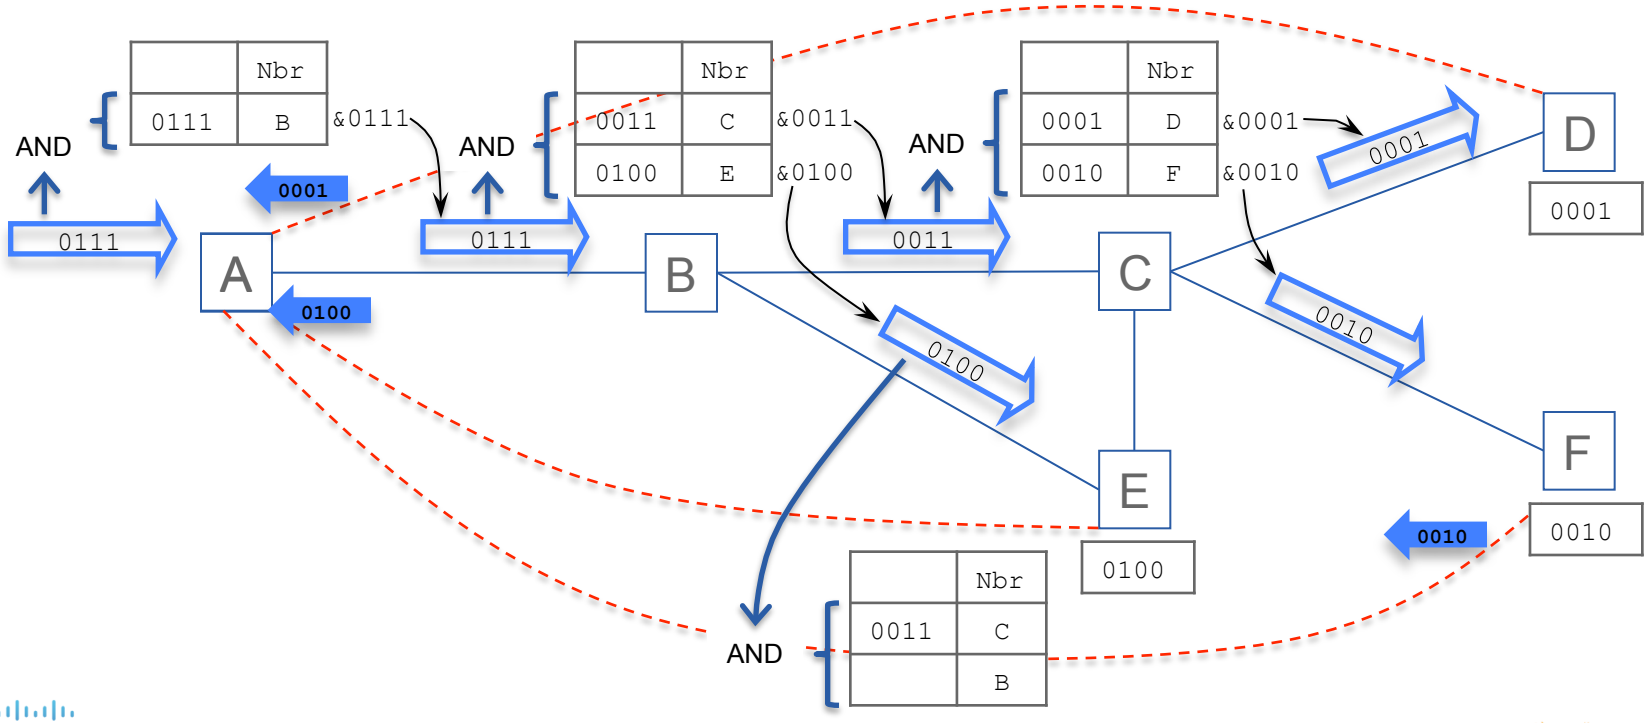

ECMP

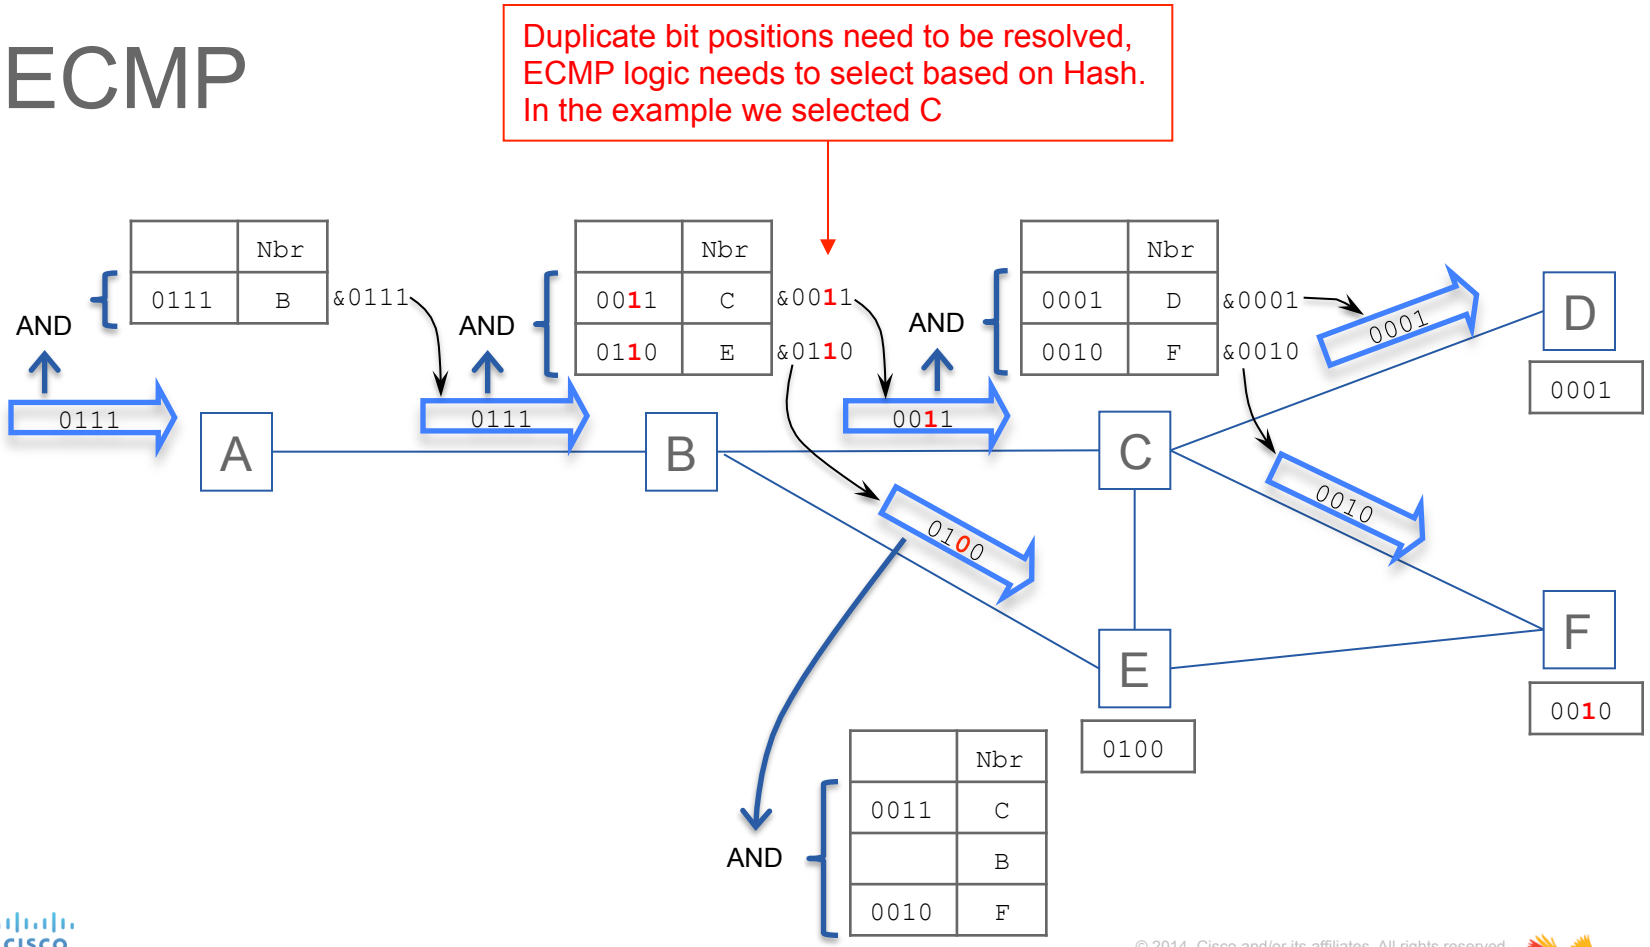

MPLS encapsulation

- We’ve analyzed the MPLS option, CRS and ASR9K platform.
- Both these platforms can do 256 bits.
- We consider 256 a good starting point.
- Other vendors confirmed they can do 256.

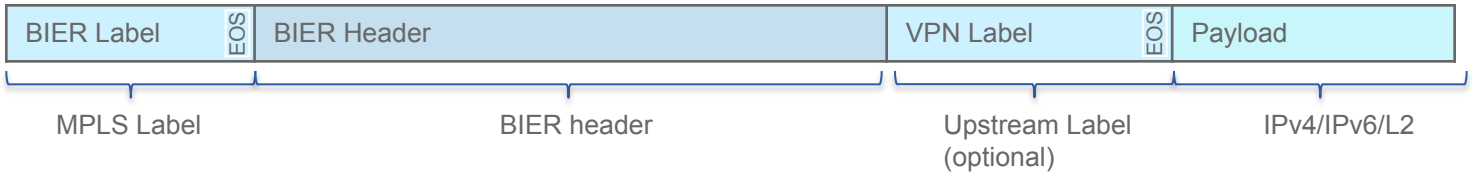

BIER Header

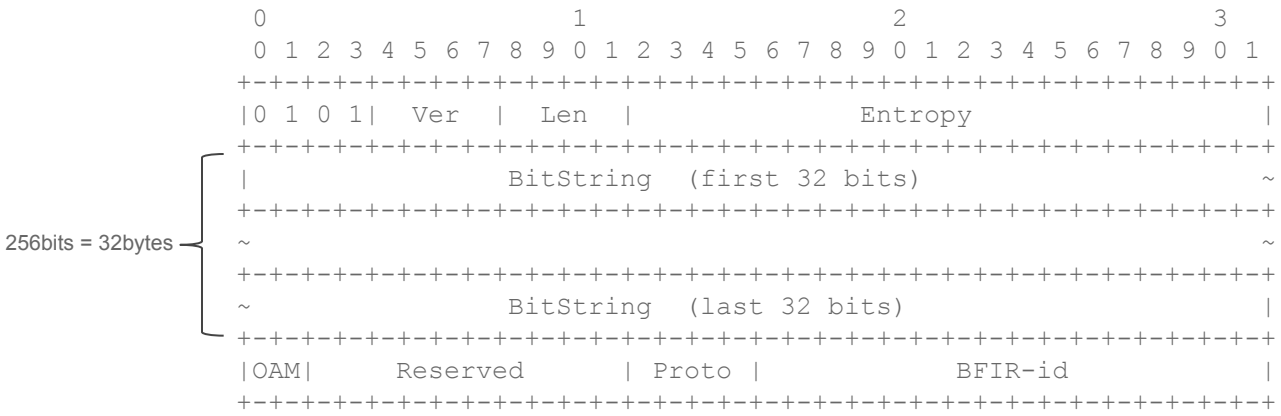

http://www.ietf.org/id/draft-ietf-bier-mpls-encapsulation-01.txt

Sets and Areas

BIER Sets

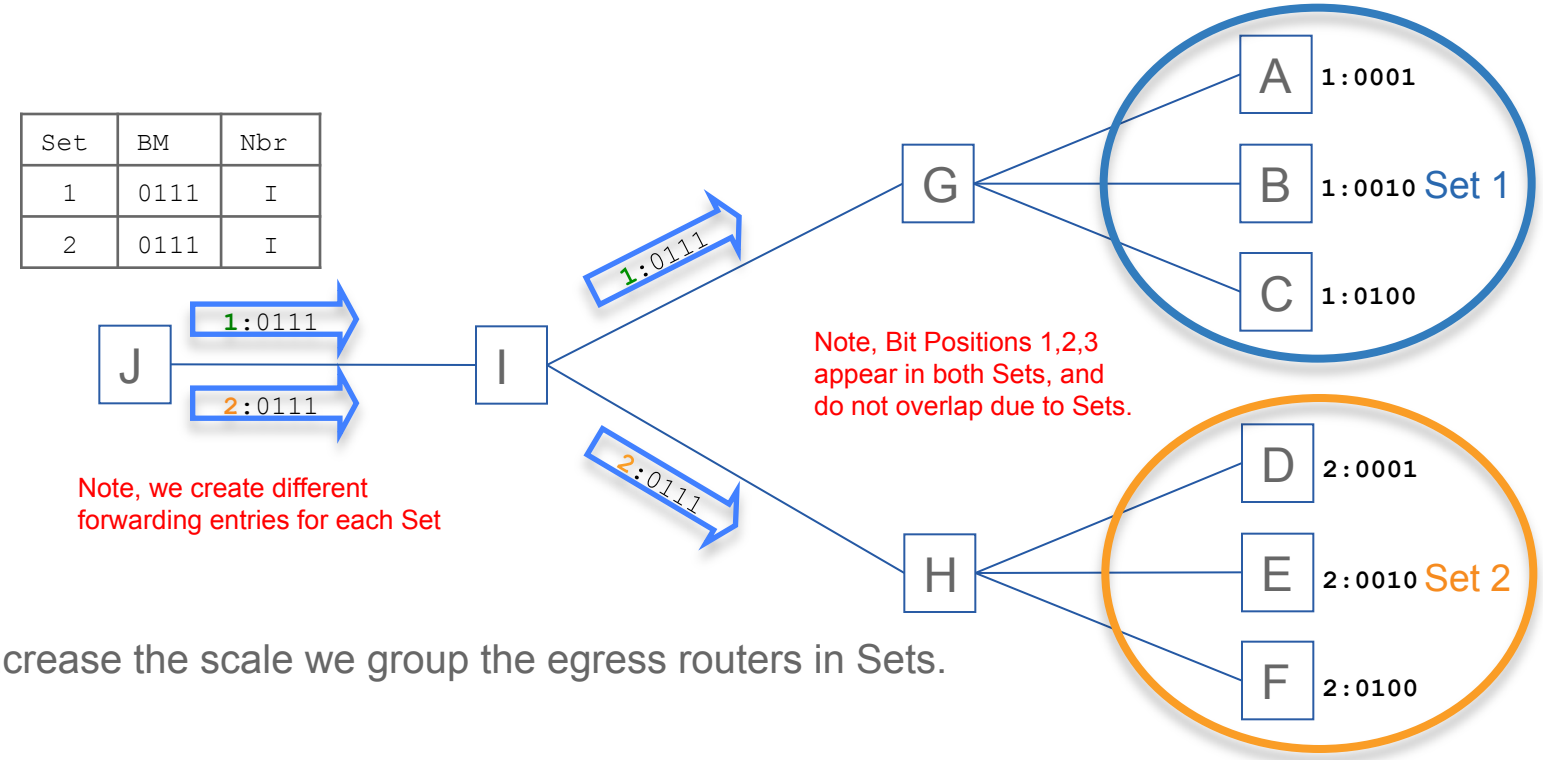

- To increase the scale we group the egress routers in Sets.

BIER Sets

Set	BM	Nbr
1	0111	I
2	0111	I

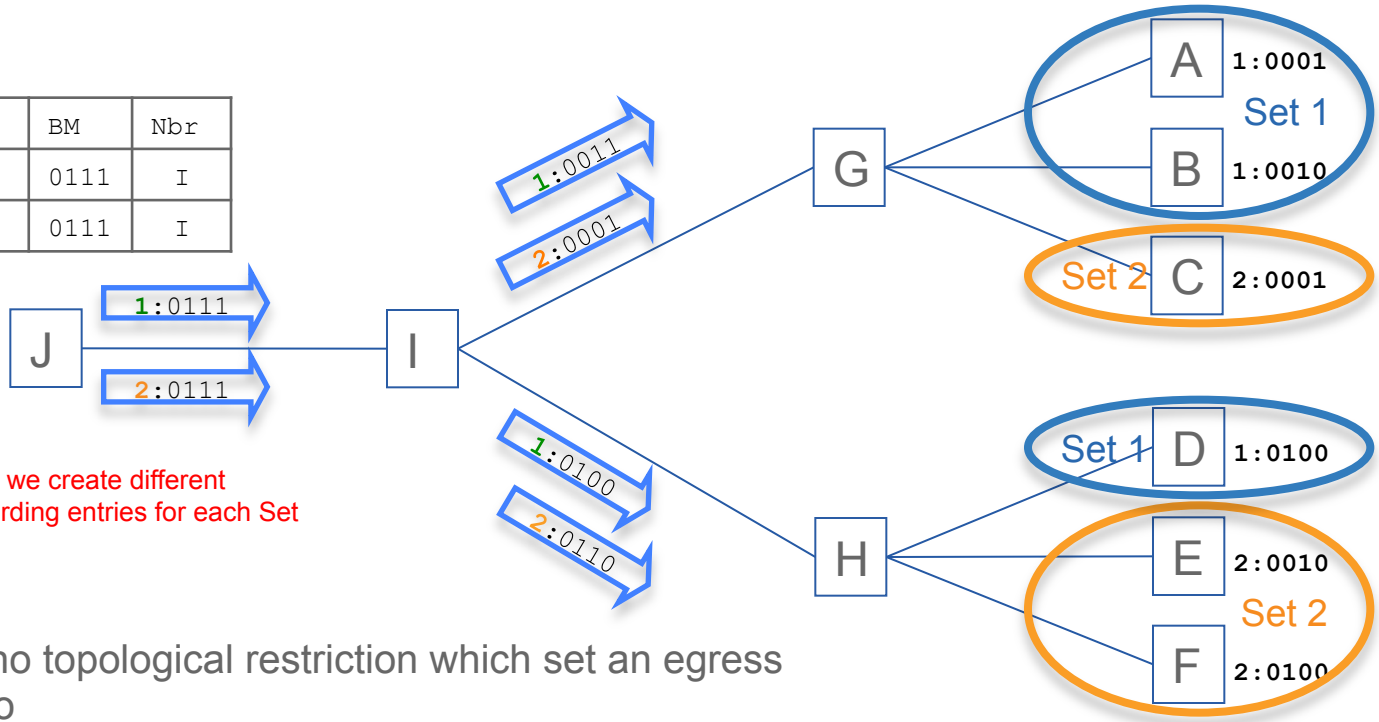

Note, we create different forwarding entries for each Set

- There is no topological restriction which set an egress belongs to

BIER Area

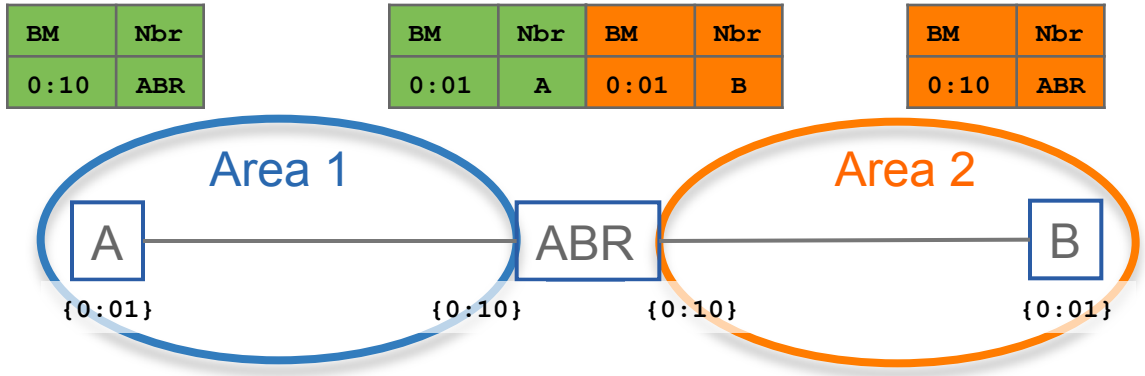

- A bit Mask only needs to be unique in its own area.
- ABR's translate Bit Masks between area's.
- Requires a IP lookup and state on the ABRs.
- This is very similar for 'Segmented Inter-AS MVPN'.

Native BIER

Native BIER

- With Native BIER there is NO PIM involved, just IGMP and BIER.
- The Source and Receiver(s) are connected to BIER router.
- There are no RP's.
- There is no equivalent of PIM modes, like sparse, ssm, bidir etc..
- We speak of 'single' sender and 'multi' sender, which is basically the same solution.
- The overlay can be BGP or SDN based.

Native BIER

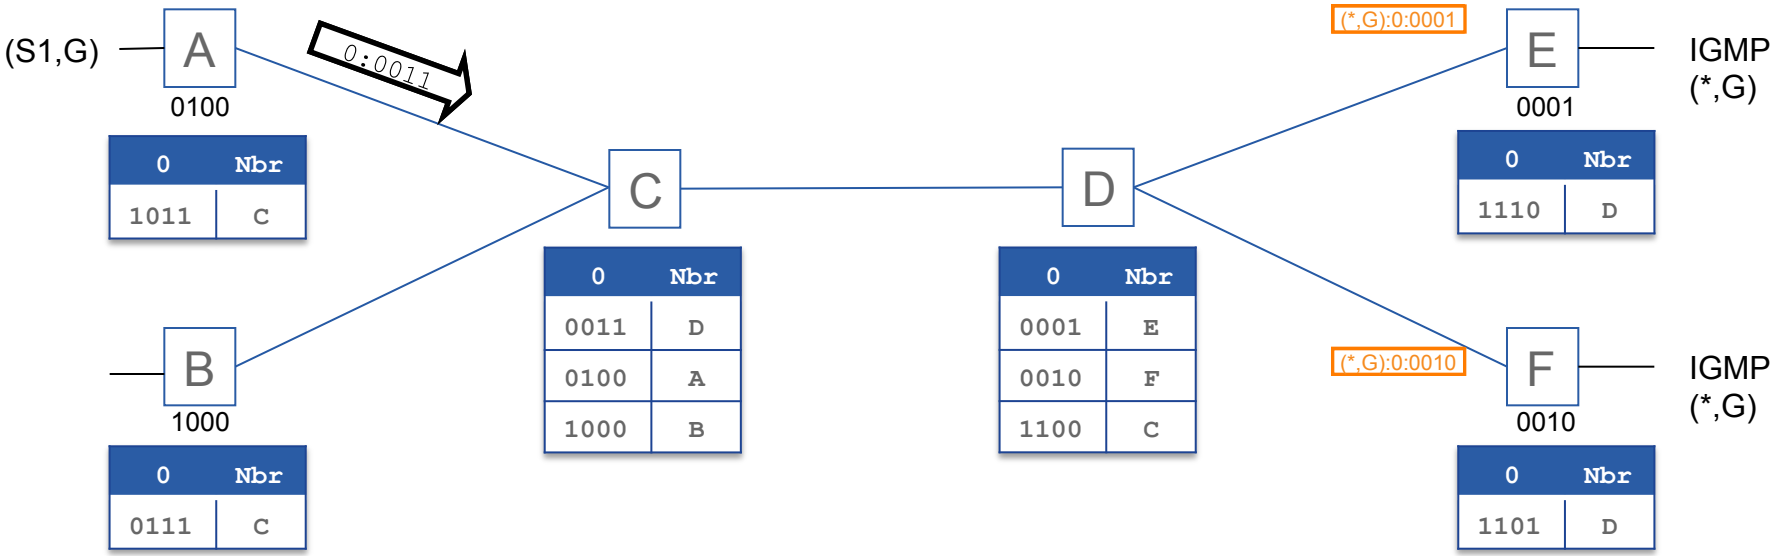

- E and F announce their Group membership via overlay to all other routers.
- A BIER router connected to the Source can immediately start sending.

Native BIER

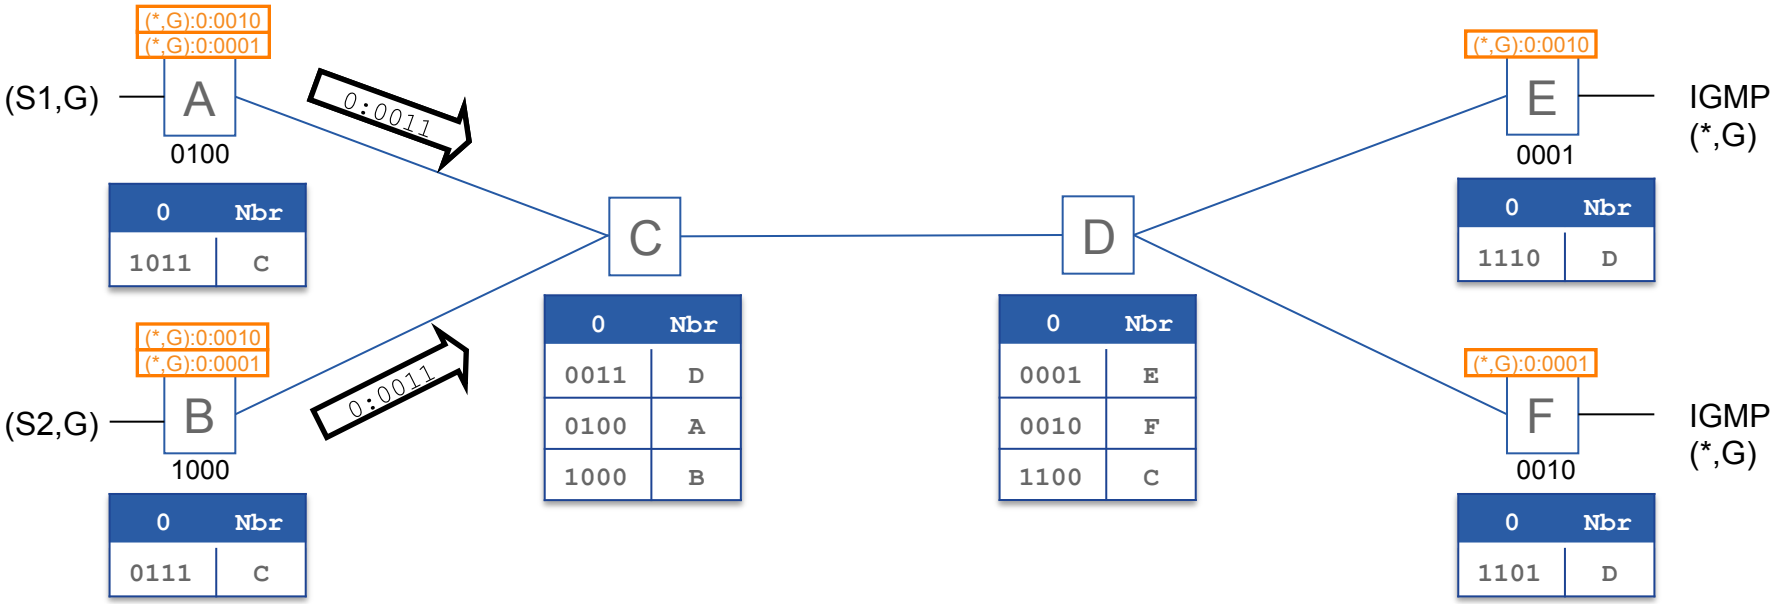

- When B leans about a new source, it can immediately start sending.

MVPN over BIER

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.

MVPN over BIER

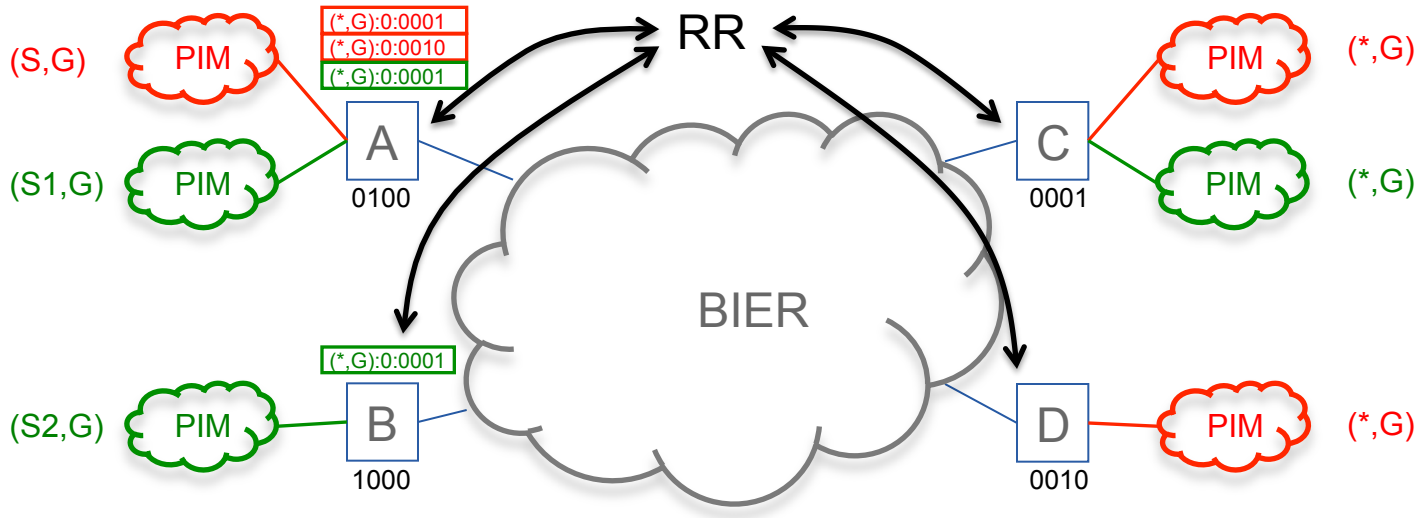

- The BGP control plane defined for MVPN can be re-used.
- Big difference, there is no Tree per VPN…!!!
- The BIER packets needs to carry Source ID and upstream VPN context label

Conclusion

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.
- Goes hand in hand with Segment Routing

Disadvantages

- The Bit String length has an upper bound and may not cover all deployment scenarios.
- Using sets to increase the number of egress routers may cause the ingress to replicate the packet multiple times.
- Using area's requires the ABR to have state.

TOMORROW starts here.