

BIER

Bit Indexed Explicit Replication

Traffic Engineering

draft-eckert-bier-te-arch-00

IETF BIER-WG Dallas 03/2015

Toerless Eckert, eckert@cisco.com

IJsbrand Wijnands, ice@cisco.com

BIER reminder

- BIER (quick reminder of terminology used)
 - Multicast replication without per-tree state.
 - BIER routers called BFR
 - Packets with BitString in header
 - Every bit (called BP – BitPosition) indicates a end-destination (BFR-id)
 - BIFT – Bit Index Forwarding Table built on every BFR
 - Every Bit Index contains IGP shortest path next-hop towards BFR-id (BFR-prefix)
 - When packet is forwarded Bits are intelligently reset to avoid loops
- 256 bits == as many bits as IPv6 address header
- Large network ? Need more tailend receivers ?
 - Send 10 packets each to 256 bits (bit-set-identifier in header)
 - Replication factor 1:256 more than enough savings for all puropses.

BIER-TE – Why

- This proposal is to support TE with the BIER mechanism

TE: “explicit path engineering” (loose, strict route through network)

Existing BIER forwarding plane can not support this

- Why “TE”

Same (or more) use cases as eg: RSVP-TE/P2MP (but without its in-network complexity)

Example : Video Contribution networks

>>50% multicast load across non-ECMP alternative paths.

Individual multicast flows > 1Gbps

- Why now ?

Define BIER-TE fwd plane + arch well enough so that we are confident that:

- a) forwarding plane is NECESSARY and SUFFICIENT for TE with BIER
- b) Control plane can be worked out now or later.
- c) Vendors can start putting BIER Fwd plane with TE option into HW.

BIER-TE – Forwarding Plane

- Every BP indicates an “adjacency”
- BIFT in each BFR ONLY populated with BPs adjacent to the BFR
- Example: p2p link BFR-A <-> BFR-B

Assign BP 33

BIFT on BFR-A: BP33 : “forward to BFR-B”

BIFT on BFR-B: BP33 : “forward to BFR-A”

BIFT on all other BFR BP 33: <not-set>

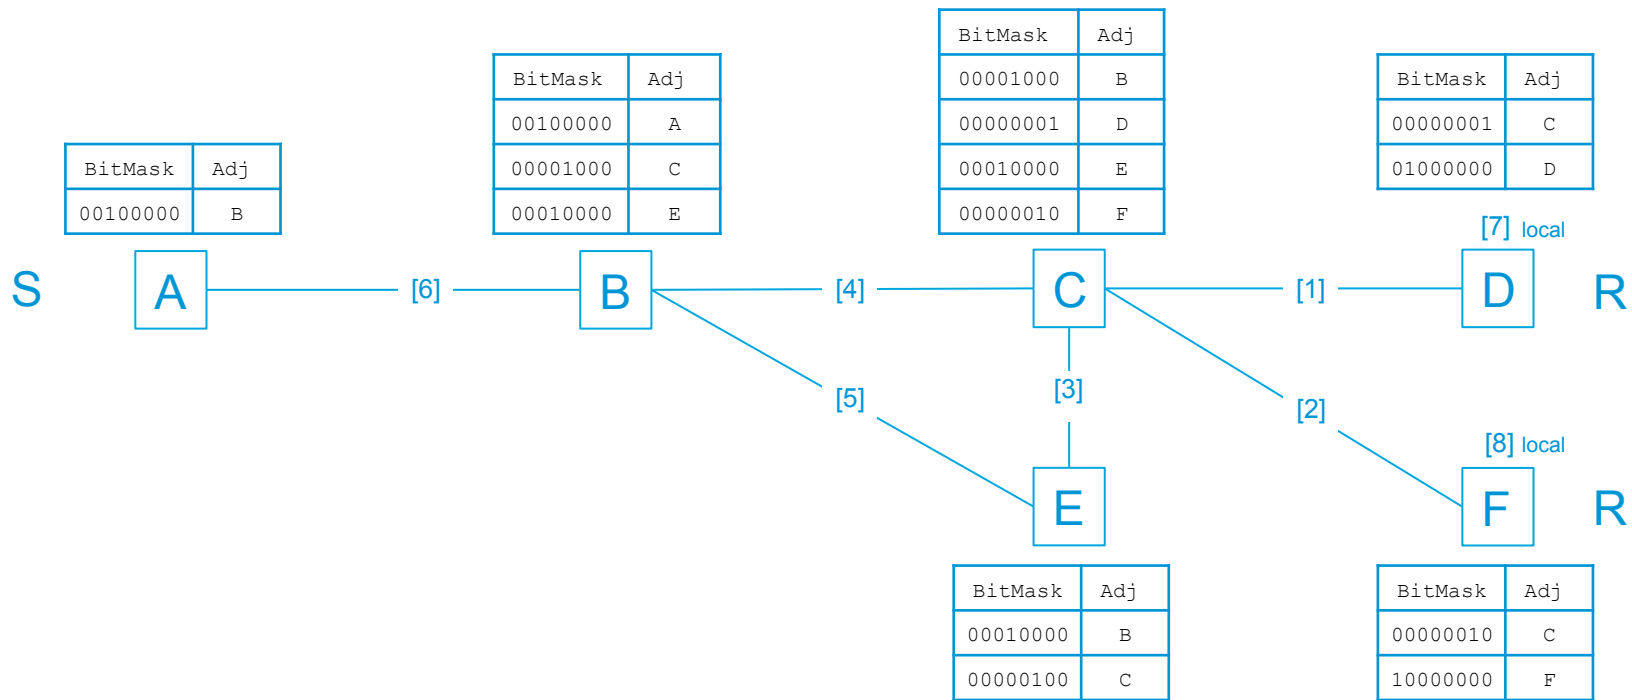
- BFR replicates BIER-TE packet to all adjacencies with BP set in packet AND BP has non-empty adjacency in BIFT

Reset BP for all adjacencies to which packet is replicated

Necessary ? Not in all cases... but extremely safe action against loops

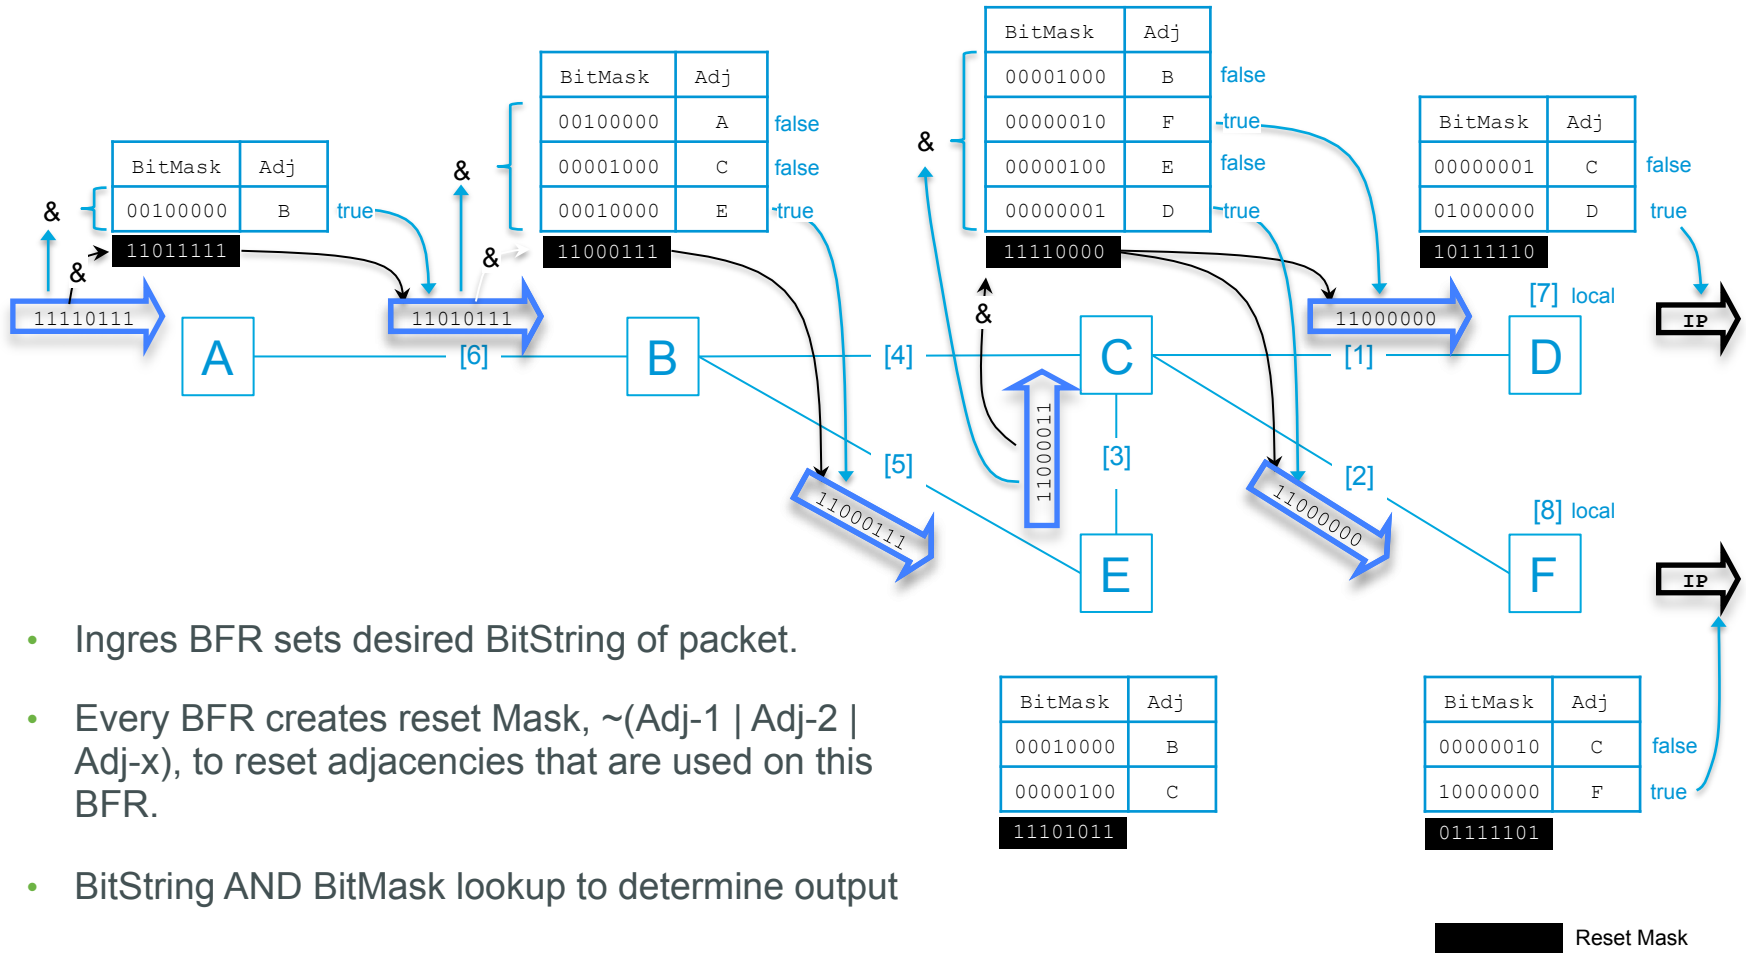
Ignore (DO NOT RESET) any other BP in packet

BIER TE Forwarding example



- All links are P2P and share BP in both directions.
- D and F have a 'local' adjacency, meaning 'For Us'.

BIER TE Forwarding example

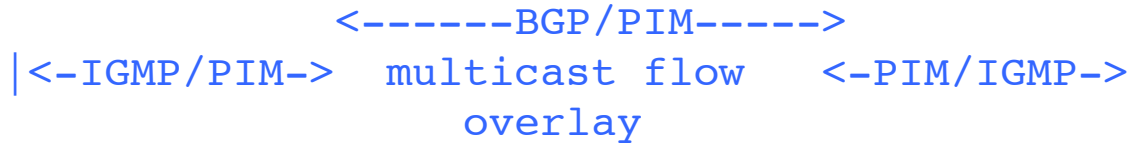


- Ingres BFR sets desired BitString of packet.
- Every BFR creates reset Mask, $\sim(\text{Adj}-1 \mid \text{Adj}-2 \mid \text{Adj}-x)$, to reset adjacencies that are used on this BFR.
- BitString AND BitMask lookup to determine output
- Update the outgoing packet using the Reset Mask

BIER-TE – Control Plane


- Control plane proposal: Central controller
- During network startup:
 - Determine topology. Calculate how to assign BPs to adjacencies
 - Installs adjacencies into BIFT of all BFR
- During network operations
 - Calculate desired traffic-engineered trees for different traffic classes
 - Get information of receivers from “multicast flow overlay”
 - Run algorithms like CSPF, Steiner, ...
 - Result of desired traffic-engineered tree is a BitString
 - Install BitString into BFIR (Ingres Router)
- Why centralized controller
 - Popular in SDN networks
 - Popular in TE solutions (PCE – Path Computation Engine)
 - Allows to minimize “in-network” complexity:
 - BIER-TE could even work without IGP in network!

BIER-TE – Summary




Same as in BIER

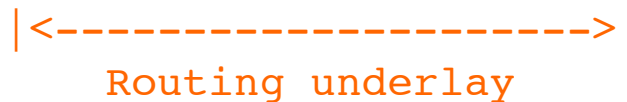
Centralized Control plane Responsible For "TE"



Src -> Rtr1 -> BFIR-----BFR-----BFER -> Rtr2 -> Rcvr



BIER-TE packets
Different semantic of bits



No IGP needed
(only for routed adjacencies)

BIER TE – large topologies

One BP per “link” requires more bits than “one bit per receiver”.

Options in -00 to reduce #BPs required:

- Single BP for whole sub-topology where “flooding” is appropriate
 - Ring, hub&spoke
 - One BP for “connected” subtopology.
 - BP only reset when packet replicated onto adjacency not part of subtopology.
- “Routed adjacency”
 - Tunnel across uninteresting parts of topology with one bit.
 - Routed link adjacency: interface-IP-address of a remote BFR
 - Routed node adjacency: loopback-IP-address of a remote BFR
- Optimize “receive “ bits for receiver-PE (as used in in BIER)



Bead inside rim creates turbulence to release flavor and aromas as beer enters mouth.

Narrowing the glass at the top retains the hop aroma and sustains the head.

Rounded shape collects aromas.

Laser etchings on bottom create bubbles for constant aroma release.

BIER

Outward turned lip delivers beer to front of tongue where sweetness (malt) is tasted.

Thinner walls and rounded shape maintain proper beer temperature longer.

Questions ?

!! NEW !!

Now engineered to your taste!