

ECMP CFM Drafting Discussion Summary January 2013

Ben Mack-Crane
(ben.mackcrane@huawei.com)

Comment #30 Resolution

ACCEPT IN PRINCIPLE. Presentation of proposal (bp-saltsidis-CFM-for-ECMP-1112-v01.pdf) and discussion in the TG meeting. Discussion covered the understanding that both VLAN MAs and PBB-TE MAs will be used in SPBM VLANs although this is not clear from reading 802.1aq. Functionality similar to PBB-TE MAs (i.e., MEPs that use specific I-SID group addresses for CCM DA) is thought to be useful (by some) in SPBM and ECMP applications. ***The next draft will include an ECMP VID MA with MEP CCM behavior that cycles through a configured set of group addresses*** (in a manner similar to the flow hash cycle currently specified for ECMP path MAs). The LTM and LBM behaviors will be the same across all MAs/MEPs/MIPs for ECMP.

Add text to 27.18 covering the use of PBB-TE MAs in SPBM VLANs.

ECMP VID MAs

- ▶ The MA *may* be associated with all I-SIDs with a particular set of endpoints (assuming “association” means that the MA tests the multicast trees assigned to those I-SIDs).
- ▶ The association is **explicit** for I-SIDs configured to be used by a MEP (i.e., the I-SID’s group address is used in CCMs).
- ▶ The association is **implicit** for I-SIDs that are not configured to be used by a MEP, but who share the same trees as some I-SID that is configured to be used by a MEP.
- ▶ There is no association for I-SIDs that are not configured to be used by a MEP and that have a tree assignment that is not used by any I-SID configured to be used by a MEP.

Note: It would be possible to configure a MEP to use an I-SID for which the MA’s endpoints are a subset of the I-SID’s endpoints without a fault being declared. In this case the MA is in some way associated with an I-SID that does not share the exact same set of endpoints.

Editor's Questions

- ▶ The MA identifier would change if MEPs are added to or deleted from the MA. Is this a desirable behavior?
- ▶ If more than one of these MAs has a MEP at a given CBP, how are these MEPs distinguished?
- ▶ Would MEP configuration be based on a list of I-SIDs, implying the use of the group addresses for those I-SIDs at that CBP?
- ▶ What behavior should be specified when an I-SID providing an address used in the MA is deleted from a CBP hosting a MEP belonging to the MA?
- ▶ What behavior should be specified when an I-SID is configured at a MEP but that I-SID is not configured at the CBP hosting another MEP in the MA?

MA Identifier

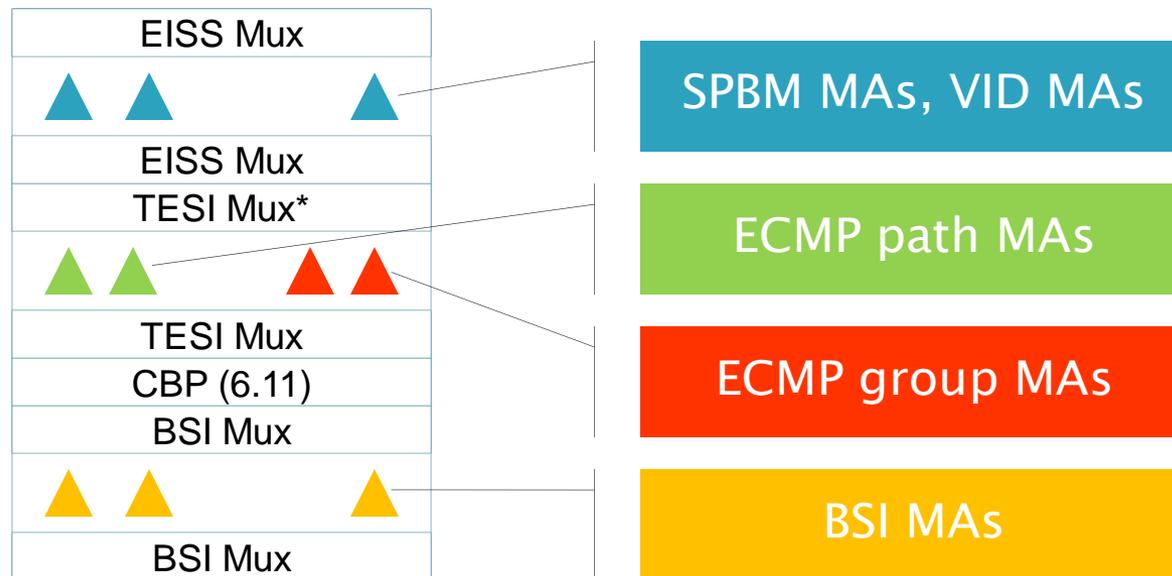
- ▶ MA to test **group address connectivity** between a specified set of CBPs
 - In the remaining slides called “**ECMP group MA**”
- ▶ The CBP set could be used as the MA identifier (list of CBP MAC addresses?)
 - If a MEP is added or removed the identifier changes
 - Delete old MA and create new MA, or
 - Change identifier on the fly.
- ▶ An I-SID could be used as the MA identifier
 - One of the I-SIDs used for monitoring, or
 - A “dummy” I-SID used only to identify the CBP set.

MEP Location

- ▶ More than one MA (MEP) may reside on a CBP
- ▶ How are frames directed to the correct MEP?
 - TESI Mux
 - BSI Mux
- ▶ In both cases the MEP is configured to send CCMs using set of I-SID group addresses
- ▶ In the **TESI Mux** case, the **group addresses associated with remote MEP's I-SIDs** must be used to select the MEP
- ▶ In the **BSI Mux** case, the set of **remote MEP's I-SIDs** must be used to select the MEP

MEP Location – TESI Mux

- ▶ TESI Mux is used for ECMP path MEPs
- ▶ Can mix group MEPs and path MEPs
- ▶ Difference in DA between MAs
- ▶ A group MEP is associated with a set of group addresses

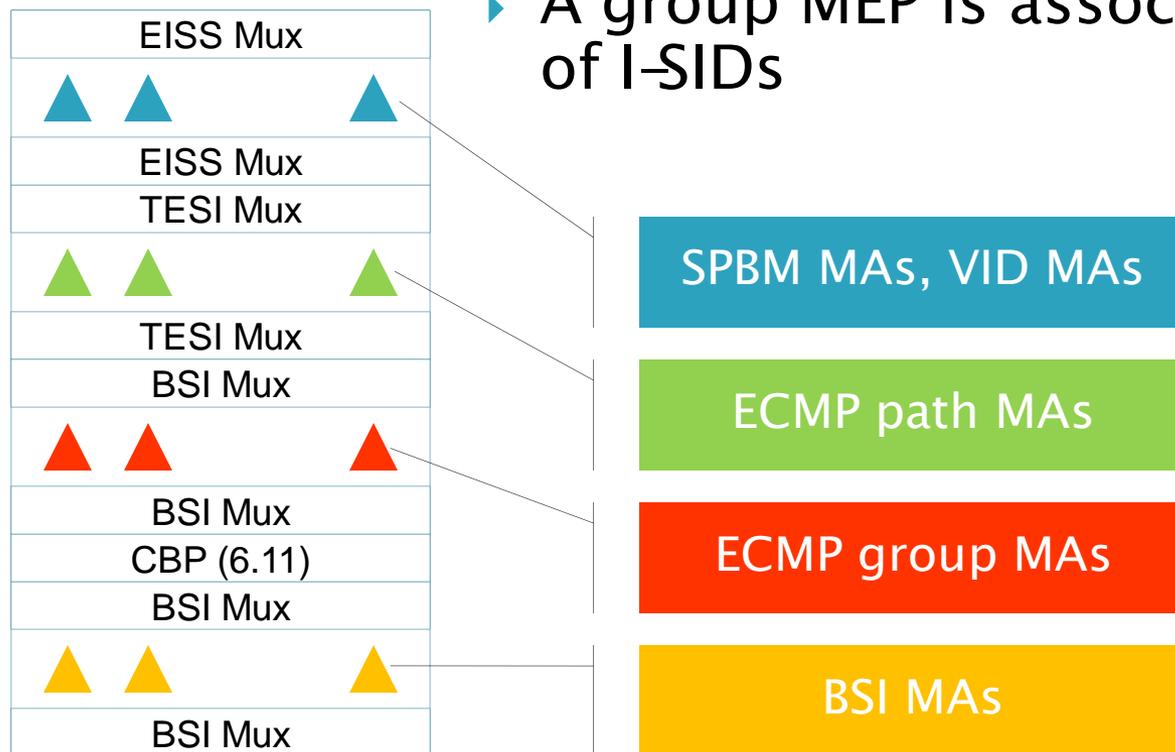


Note: Only the SPBM MA can reach MIPs (i.e., perform LB and LT) since it is at the lowest MD level on the CPB.

*Note: Figure does not show proper Mux nesting

MEP Location – BSI Mux

- ▶ BSI Mux is used for ECMP path MEPs
- ▶ A group MEP is associated with a set of I-SIDs



Notes: Only single I-SID per MEP w/BSI Mux;
No CPB address translation – so this may not work,
but this may.

MEP Configuration Details

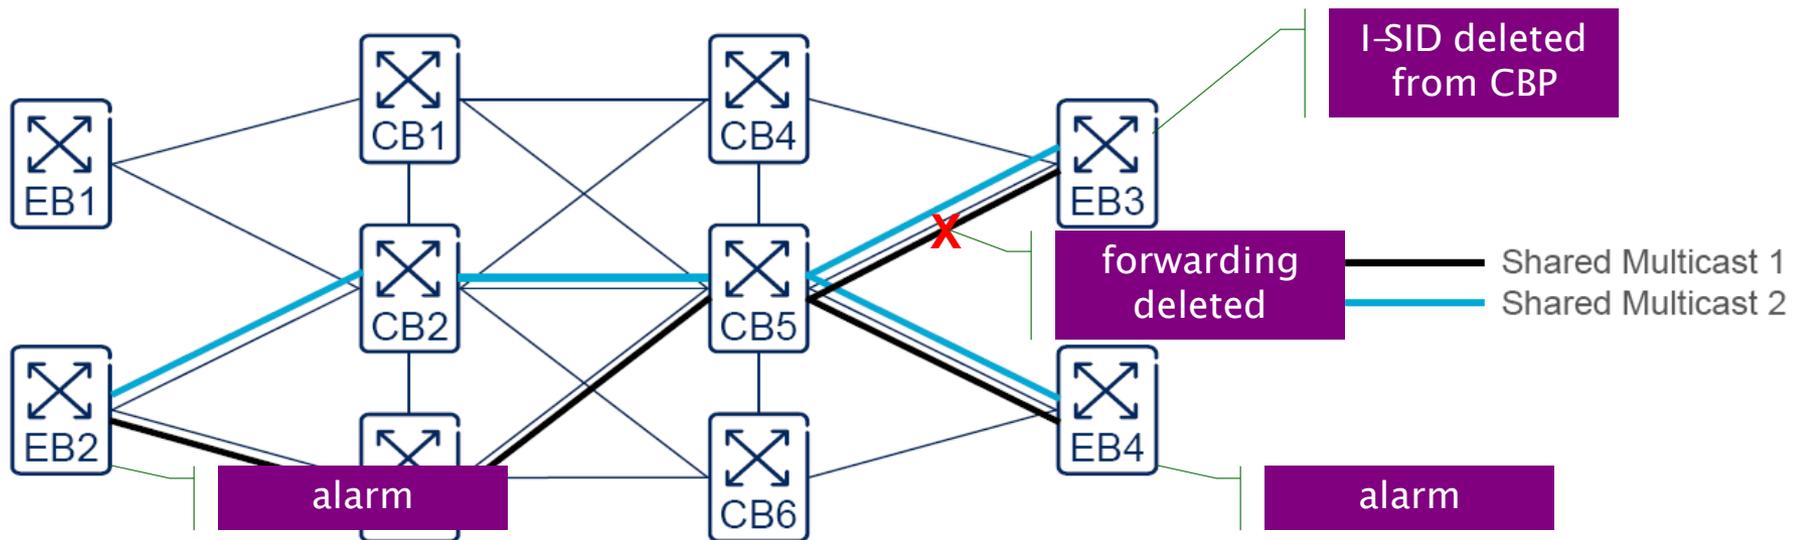
- ▶ Each MEP in an ECMP group MA is configured with a **list of I-SIDs** it will use for CCM DAs
- ▶ In a **TESI Mux**, the MEP must be associated with specific 3-tuples $\langle \text{DA}, \text{SA}, \text{VID} \rangle$
- ▶ Therefore, an ECMP group MEP must be configured with the 3-tuples from all the remote MEPs in the MA
- ▶ For each **remote MEP x** , and each **I-SID i** configured in MEP x one must determine the group address used by the I-SID at that CBP and record a **3-tuple $\langle \text{DA}_i, \text{SA}_x, \text{VID} \rangle$** in the TESI Mux for the local MEP

Note: If an I-SID is configured not to send multicast frames at a CBP, then that I-SID cannot be used by the MEP.

Configuration Consistency

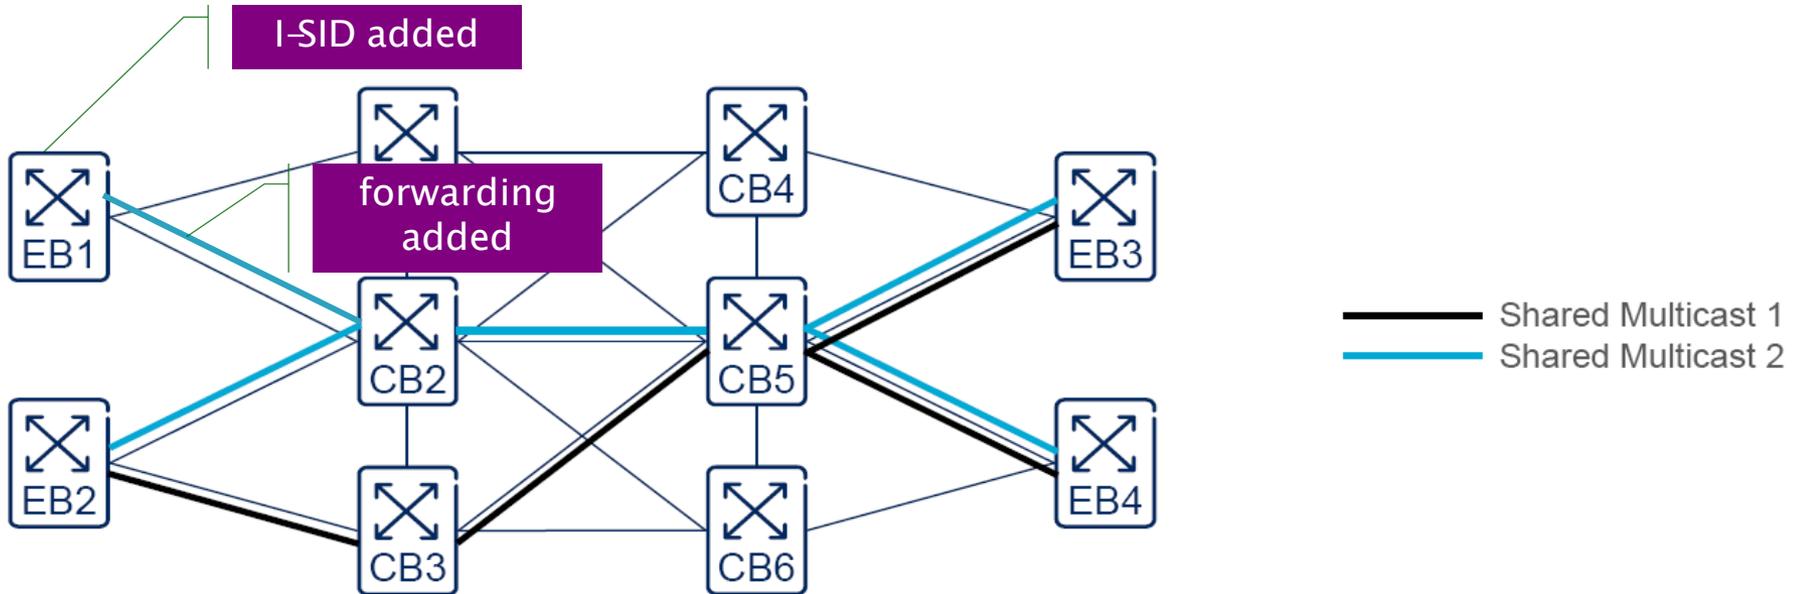
- ▶ An I-SID used in an ECMP group MA MUST be configured on (at least) the CBPs in the MA.
- ▶ It was noted that automation could ensure consistent configuration of MEPs in case of changes in I-SID connectivity
- ▶ Specifying automation is a significant task
- ▶ As noted in San Antonio, automation of MEP configuration is not in scope for Qbp
- ▶ Qbp will only address the MEP configuration required and any consistency rules that must be maintained by an operator or controller

Deleting an I-SID Endpoint



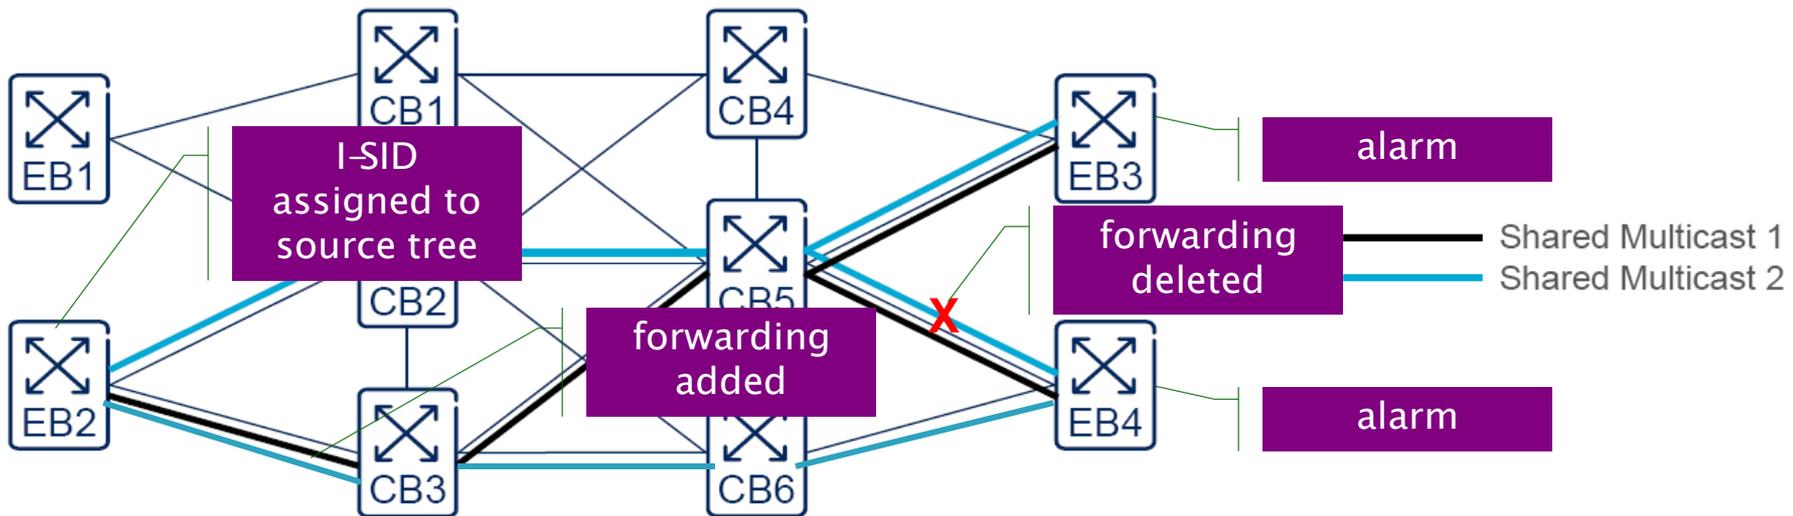
- ▶ If an I-SID is deleted from a CBP the forwarding state for that I-SID's group address at that CBP will be deleted.
- ▶ If that I-SID is configured in an ECMP group MEP, the path for CCM using the deleted address may be broken.
- ▶ Therefore, **before deleting an I-SID from a CBP, one should remove the I-SID from any ECMP group MA.**

Adding an I-SID Endpoint



- ▶ If an I-SID is added to a CBP, forwarding for that I-SID's group addresses to that CBP will be created.
- ▶ If that I-SID is configured in an ECMP group MEP, CCMs will arrive at the new CBP and be discarded.
- ▶ No defect will be detected, however this situation does not conform to the strict consistency rule originally proposed.

Changing an I-SID Tree



- ▶ If an I-SID is assigned a new multicast tree type, the forwarding state for that I-SID's old group address will be deleted and a new address added.
- ▶ If that I-SID is configured in an ECMP group MEP, the expected address at receiving MEPs will be wrong.
- ▶ Therefore, **before changing an I-SID's tree type at a CBP, one must add the new expected address to the other MEPs in the MA.**

ECMP group MA Maintenance

- ▶ To maintain the correspondence between I-SID endpoints and ECMP group MA endpoints
 - If an I-SID endpoint is added
 - the I-SID should be removed from any ECMP group MA, and
 - A new ECMP group MA created for that I-SID, or
 - Add the I-SID to an* ECMP group MA with matching endpoints, or
 - No longer use the I-SID in an ECMP group MA.
 - If an I-SID endpoint is to be deleted
 - The I-SID should first be deleted from any ECMP group MA
 - Delete the I-SID from the CBP, and
 - Create a new ECMP group MA for that I-SID, or
 - Add the I-SID to an ECMP group MA with matching endpoints, or
 - No longer use the I-SID in an ECMP group MA.
 - If an I-SID's multicast tree type changes
 - Add the new expected address before changing
 - Remove the old expected address after changing

*Can more than one ECMP group MA have the same set of endpoints?

ECMP group MA Maintenance

- ▶ Adding an ECMP group MA MEP
 - For each I-SID used in the MA
 - Add an I-SID endpoint at the new CBP, or
 - Remove the I-SID from the MA
- ▶ Deleting an ECMP group MA MEP
 - For each I-SID used in the MA
 - Delete the I-SID at the CBP where the MEP is deleted, or
 - Tolerate the looser consistency rule.

BSI MAs

- ▶ BSI CFM is already defined
- ▶ One I-SID per MA
- ▶ Capable of monitoring the same group address trees as ECMP group MAs
- ▶ **No coordination required** with backbone MAs when making changes to BSI endpoints
- ▶ Select a set of BSIs to monitor the desired multicast trees and enable BSI CCMs
- ▶ If a selected BSI's endpoints change, reevaluate the selection, if necessary

Choices

- ▶ Do we add ECMP group MAs to Qbp?
 - How are ECMP Group MAs identified?
 - Single I-SID per MA* or multiple I-SIDs per MA?
- ▶ Do we describe how BSI MAs can be used to monitor ECMP multicast trees?

*Reduces complexity of MA maintenance by removing constraints related to I-SID grouping.

Use of PBB-TE MAs with SPBM

- ▶ From Janos:
 - aq-farkas-CFM-in-802.1aq-0908.pdf
 - aq-farkas-proposal-for-CFM-in-SPB.pdf
- ▶ For SPBM MAs the SA in LBR and LTR frames should be set according to the same logic used for PBB-TE
 - not currently specified in the 802.1Qbp draft
- ▶ Does not describe how MAs are configured
- ▶ Is anyone convinced that we need to specify PBB-TE MAs to be used with SPBM VLANs?