



PBB-TE (802.1Qay) CFM

Dinesh Mohan

mohand@nortel.com

Sep 4-7, 2007

802.1 Stockholm Interim



Agenda

- Recap from last meeting
- CFM requirements in PBB-TE
- CFM enhancements needed in PBB-TE
- Enhancement Solutions



Recap

- In Jul'07 San Francisco plenary, [ay-mohan-cfm-0707-v01.pdf](#) was presented
- Lot of discussion triggered on PBB-TE CFM requirements and assumptions made
 - CCM related enhancement was generally agreed
 - LBM/LBR and LTM/LTR related discussion did not conclude due to time constraints
 - Main aspect whether ESPs are co-routed or not in either direction
 - Presentation assumed co-routed ESPs
 - Whether B-VID in either direction is same or can be different
 - Presentation assumed general case where B-VIDs could be different
- Current presentation is a follow-up to July's presentation

PBB-TE Overview



- PBB-TE is based on provisioned Ethernet Switched Paths (ESPs) in PBBN
 - ESPs can be provisioned via management or control plane
 - B-VID is one of the PBB-TE assigned B-VID
- Presentation Update: Draft .1Qay/0.0 assumes that ESP is a bidirectional path
 - Current definition: *ESP is a provisioned path between two VIPs which extends over a PBBN*
 - VIP should really be CBP
 - [ay-mohan-cfm-0707-v01.pdf](#) assumed a bidirectional path with identified being <B-DA, B-SA, B-VID1, B-VID2>; VID1 and VID2 could be the same
 - Appendix used a different interpretation of ESP as being unidirectional identified by <B-VID, B-DA> tuple
 - Bidirectional path view is accurate with two directions identified by <VID,DA>; this presentation will be updated accordingly – however does not impact the contents
- **Co-routed** unidirectional ESPs constitute bidirectional PBB-TE path
 - Co-routing ESPs preserves bidirectional Ethernet behavior
 - Co-routing ESPs simplifies diagnostics
 - Co-routing ESPs preserves application of CFM mechanisms
- This presentation assumes co-routed ESPs
 - PBB-TE is 802.1 activity which is expected to operate over MACs and physical layers specified by other 802 groups, mostly 802.3. To the best of knowledge, 802.3 does not support one way service. Further, there are a number of assumptions in 802.3 to preclude one way service. Historically, one way service has been deemed to be dangerous to the network.



PBB-TE CFM Requirements

- PBB-TE path continuity fault detection
 - **R1**: Between two end-points of PBB-TE path, proactively detect continuity break (e.g. sub-50ms failure detection)
- PBB-TE path continuity fault verification
 - **R2**: Verify on-demand connectivity between two end-points of PBB-TE path
 - **R3**: Verify on-demand connectivity between an end-point and an intermediate point of PBB-TE path
- PBB-TE path continuity fault localization
 - **R4**: Isolate fault location when continuity between two end-points of PBB-TE path fails
 - **R5**: Determine the route of a PBB-TE path



R1: Continuity Fault Detection between MEPs

- Between two end-points of PBB-TE path, proactively detect continuity break (e.g. sub-50ms failure detection)
 - Requirement is applicable in both EMS/NMS dependent environments as well as those without EMS/NMS
 - Continuity break can be carried out in unidirectional manner therefore applicable to individual ESPs
- Enhancement: **Unicast CCM**
 - Unicast address is same as B-DA in direction of ESP
 - This is such that forwarding along the path is based on same <B-DA, B-VID> tuple as any data path frame as required by PBB-TE ESP
 - Unicast CCMs are already supported in Y.1731 and are not precluded in 802.1ag
 - For explicit support, update will be needed to text from .1ag/D8.1 c3.2, c8.13.11, c18 etc. where multicast is mentioned currently

R2: Continuity Fault Verification between MEPs



- Verify on-demand connectivity between two end-points of PBB-TE path
 - Requirement is applicable to both EMS/NMS dependent environments as well as those without EMS/NMS
 - Loopback mechanism defined in 802.1ag can be applied which means both ESPs are used
 - One ESP for LBM and other ESP for corresponding LBR
- Enhancement
 - **None**
 - If same B-VID is used in either direction of PBB-TE path
 - **Change B-VID in LBR**
 - If different B-VIDs are used in either direction

R2: Continuity Fault Verification between MEPs (cont'd)



- Verify on-demand connectivity between two end-points of PBB-TE path when different B-VIDs used in either direction of PBB-TE path
- Enhancement: **Change B-VID in LBR**
- Different solution options:
 - Option 1: Carry a TLV with reverse VID in LBM which is used by loopback point for VID value in LBR
 - Advantage: Makes LBM sink stateless and processing in LBM sink generic, i.e. if a specific TLV present, use its value for LBR; however makes LBM source add TLV for PBB-TE ESP
 - Option 2: Since PBB-TE ESP MEP maintains association between forward and reverse VIDs, have loopback point perform this VID change
 - Advantage: Make LBM source point generic; however makes LBM sink stateful
 - If primary VID of MEP in one direction is configured with that direction's VID, this functionality already exists today
- Recommended solution (in consideration of R3): Carry TLV with reverse VLAN

R3: Continuity Fault Verification between MEP & MIP

- Verify on-demand connectivity between an end-point and an intermediate point of PBB-TE path
 - Requirement is mostly applicable to environments which are not dependent on presence of EMS/NMS – since application is to diagnose intermediate problems which will generally be detected in EMS/NMS
 - Loopback mechanism defined in 802.1ag can be applied which means both ESPs are used
 - One ESP for LBM and other ESP for corresponding LBR
 - Not possible if ESPs are not co-routed
- Enhancement: **Selective intercept of LBMs at MIP**
 - MIP should be able to intercept LBMs intended for it and should be able to ignore LBMs not intended for it
 - LBM frames, intended for MIPs should have DA corresponding to ESP B-DA
 - If DA in LBM is MIP's MAC, MIP MAC may not be provisioned in filtering databases associated with ESP B-VIDs, meaning that LBM frame may be discarded since flooding is not allowed
- Enhancement: **Change B-VID in LBR**
 - Same requirement as R2

R3: Continuity Fault Verification between MEP & MIP (cont'd)



- Verify on-demand connectivity between an end-point and an intermediate point of PBB-TE path
- Enhancement: **Selective intercept of LBMs at MIP**
- Among the three options discussed during last meeting
 - New EtherType and New OpCode options ruled out due to RAC considerations and inability to provide selective intercept
 - Option to carry a TLV seen to be viable
 - A TLV to be used as first TLV to allow deterministic inspection at intermediate MIPs
 - As per current format, this would imply looking at 10-15 octets following OpCode
 - The first field in Value of TLV is MIP identifier i.e. MAC address, which allows MIPs to selectively intercept CFM frames intended for it
- Enhancement: **Change B-VID in LBR**
- Different solution options:
 - Option 1: Carry a TLV with reverse VID in LBM as also indicated in R2 option
 - Advantage: Makes LBM sink stateless and processing in LBM sink generic, i.e. if a specific TLV present, use its value for LBR; however makes LBM source add TLV for PBB-TE ESP
 - Option 2: Expect MIP to maintain association between the two direction VIDs
 - Configuring the primary VID of MHF via MA configuration is possible however different primary VID for each MHF associated with a MIP is not possible which will be needed for this option
- Recommended solution: Carry first TLV with Loopback point's MAC address and reverse VLAN



R4: Continuity Fault Localization

- Isolate fault location when continuity between two end-points of PBB-TE path fails
 - Requirement is mostly applicable to environments which are not dependent on presence of EMS/NMS – since application is to diagnose intermediate problems which will generally be detected in EMS/NMS
 - Linktrace mechanism defined in 802.1ag can be applied; means both ESPs are used
 - One ESP for LTM and other ESP for corresponding LTR
 - Not possible if ESPs are not co-routed
- Enhancement: **Change B-VID in LTR**
 - If DA in LTM is a multicast MAC as per Table 8-10/802.1ag/D8.1:
 - Static entry for this group MAC address will need to be added in all devices a priori
 - since VID can be reused across different ESPs, LTM would not be bounded to only PBB-TE ESP path
 - since target MAC may not be provisioned in filtering databases associated with ESP VIDs, MIP would have no means to determine whether or not they are in the path of ESP for that VID
 - LTM frames, should have DA corresponding to ESP B-DA
- Recommended solution: Carry TLV with reverse VLAN



R5: PBB-TE Trunk Route Determination

- Determine route of a PBB-TE path
 - Requirement is mostly applicable to environments which are not dependent on presence of EMS/NMS – since EMS/NMS dependent environment would already have a view of PBB-TE path route
 - Linktrace mechanism defined in 802.1ag can be applied; means both ESPs are used
 - One ESP for LTM and other ESP for corresponding LTR
 - Not possible if ESPs are not co-routed
- Enhancement: **Change B-VID in LTR**
 - Same as R4 (last slide)



Proposed TLV

	Octets
Type = 9	1
Length	2-3
MIP MAC	4-9
Reverse VID	10-11

PBB-TE ESP TLV

- Applicable to R2, R3, R4, and R5
 - R2 (LBM/LBR between MEPs) requires TLV to carry **Reverse VID**; **MIP MAC** is don't care
 - R3 (LBM/LBR between MEP/MIP) requires TLV to be first TLV with both **MIP MAC** and **Reverse VID**
 - R4 & R5 (LTM/LTR) require TLV to carry Reverse VID; **MIP MAC** is don't care



Backup slides



➤ Option 1: New EtherType

➤ Pros:

- facilitates datapath to differentiate between CFM frames for MEPs & MIPs

➤ Cons:

- means duplicate EtherTypes for same functionality – bad!
- not a requirement for PBB-TE MEPs e.g. CCM, LBMs etc.
- Every MIP along ESP path before destination will process frame

➤ Option 2: New OpCode

➤ Pros:

- facilitates datapath to differentiate between CFM frames for MEPs & MIPs

➤ Cons:

- means duplicate OpCodes for same functionality – bad!
- not a requirement for PBB-TE MEPs e.g. CCM, LBMs etc.
- Every MIP along ESP path before destination will process frame



➤ Option 3: New TLV

➤ Pros:

- facilitates datapath to **selectively** differentiate between CFM frames for MEPs & MIPs
- Does not lead to duplication of EtherType or OpCode
- Consistent with current 802.1ag/Y.1731 design

➤ Cons:

- Requires packet inspection at MIPs datapath to support efficient usage
- Not a requirement for PBB-TE MEPs CCM

➤ Proposed Solution: Use Option 3 (see subsequent slides)

Enhancement#3 – MIP LBM



➤ Enhancement#3:

- PBB-TE ESP MIPs should be able to intercept LBMs intended for it
- PBB-TE ESP MIPs should be able to ignore LBMs not intended for it

➤ Proposed solution:

- A TLV to be used as first TLV to allow deterministic inspection at intermediate MIPs
 - As per current format, this would imply looking at 10-15 octets following OpCode
- The first field in Value of TLV is MIP identifier i.e. MAC address, which allows MIPs to selectively intercept CFM frames intended for it