Notes on ECMP CFM
November 2012

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

Fault Hypothesis (very rough)

- Faults that affect all traffic
  - Link failure
  - Node failure
- Faults that affect specific flow(s)
  - Memory error
  - FDB entry error

Why not test all paths?

- Complexity in determining cover set
  - Which (minimal?) set of addresses covers all paths
  - How many peers does each MEP expect (or how many MEPs/MAs are required)?
  - How does this change with topology?
- Faults that affect all traffic are recognized and routed around (assuming some connectivity remains)
  - So path test CCMs are rerouted over remaining connectivity
  - I.e. they do not actually test specific paths
- Faults to specific CCM flows are expected to be extremely rare
  - And therefore not particularly useful
- Monitoring all FDB state (CCM for every address) does not scale
  - And has diminishing usefulness
- Therefore, decided one endpoint reachability monitor was sufficient
  - Default I-SID group address state automatically installed – use this!
  - Can instantiate MA to monitor entire B-VID or endpoints for a given service

Why allow arbitrary LTM DA?

- In xSTP controlled VLANs, LTM uses a reserved group DA
  - Floods through Bridges that do not have MIPs
  - Only one path in spanning tree, so will find MIPs further on that can provide LTRs
- In SPBM
  - No flooding; and No forwarding state for LTM reserved addresses
  - **Could** use source specific group address for default I-SID; **however**
  - All FDB state is computed and installed (not learned)
  - Increased possibility of errant FDB entries that do not follow "expected" path
  - If a Bridge without MIPs is traversed by LTM and the default address is forwarded differently from target address (**a bug!**) the trace **will not find** the path of interest further on…
    - May find path that looks good, but is not the real path
    - May find no path, but there really is a path somewhere else

Better to **use target address as the DA** so **real path is traced** accurately through Bridges with no MIPs!

What has this to do with ECMP?

- If we do not need to monitor all paths;
- If LTM should use target address as DA;
- Then the “SPBM MA” and “ECMP VID MA” function in the same way!
  - Of course the F-TAG information must be included when using flow filtering
- The current draft is based on these decisions
  - See next slide for details (presented in April & May)

802.1Qbp CFM – in one slide

CCM VLAN	DA is SPBM default I-SID SPsourceID+00-00-FF	Take advantage of installed forwarding state for SPBM default I-SID All to all CCMs, provision MAID and expected MEP IDs Tests endpoint reachability, not all paths
CCM path	DA is individual address of CBP Cycle through Flow Hash values	Tests multiple paths between two points; use TE-SID to identify MA MEPs located in TESI multiplexer ; <SA, DA, VID> selection Send each Flow Hash 4 times to cause RDI in case of path fault Correlate RDI with Flow Hash via cycle location, sending rate, path delay
LBM	DA is any individual/group address Use PBB-TE MIP TLV to target MIP	Use same MIP datapath as for PBB-TE
LBR	DA is LBM SA	No change here from VLAN CFM case
LTM	DA is any individual/group address Flow Hash for individual address	Same rules as PBB-TE, allowing multiple Egress ports Use flow hash in FDB lookup, if required
LTR	DA is Original MAC Address from LTM PDU	No change here from VLAN CFM case