ERICSSON

CFM FOR ECMP

Panagiotis Saltsidis

CFM

- Provides capabilities for detecting, verifying and isolating **connectivity** failures end-to-end through an Ethernet bridged network
- The following mechanisms are supported:
  - Continuity Check (CC)
  - Loopback (LB)
  - Linktrace (LT)
- Designed for Bridged Networks that may be operated by multiple independent organizations, each with restricted management access to each other's equipment

MAINTENANCE DOMAINS & ASSOCIATIONS

- A Maintenance Domain (MD) is the network or the part of the network for which faults in connectivity are to be managed.
- An MD can provide service instances to an enclosing MD, or utilize service instances provided by an enclosed MD. An explicit Maintenance Domain Level (MD Level) is included in CFM PDUs to indicate the nesting relationships among MDs
- A Maintenance Association (MA) is set to monitor the **connectivity** of a specific service instance within an MD
- The configuration of a Maintenance Association (MA) requires a parameter (or a list of parameters) that associates the MA with the monitored service and in practice enables control of the CFM PDU structure in order to enable “fate sharing” with the monitored data.

VID SEMANTICS

- The VID parameter in IEEE802.1Q has overloaded semantics. A C-VID or S-VID is used
  1. As a control protocol identifier
  2. As a VLAN connectivity identifier
  3. As an active topology selector
  4. As an end-to-end service identifier

BRIDGING EVOLUTION – IEEE802.1ah

- Newer versions of the standard are modifying the original set
- A B-VID is used:
  1. As a control protocol identifier
  2. As a VLAN connectivity identifier
  3. As an active topology selector

  But **NOT** as an end-to-end service identifier.

  I-SIDs are used for this.

BRIDGING EVOLUTION – IEEE802.1aq

- An SPBM-VID (subset of a B-VID) is used:
  1. As a control protocol identifier
  2. As a VLAN connectivity identifier
  3. But **NOT** as an active topology selector.

     An SPBM-VID alone cannot select an active topology.

     SPBM-VID + MAC SA is used for this
  4. And (carrying the properties of a B-VID) **NOT** as an end-to-end service identifier.

     I-SIDs are used for this.

BRIDGING EVOLUTION – IEEE802.1Qbp

- An ECMP-VID (subset of an SPBM-VID) is used:
  1. As a control protocol identifier
  2. But **NOT** as VLAN connectivity identifier.

     The connectivity associated with an ECMP-VID does not map to a VLAN connectivity.

     There is currently no identifier used for this
  3. And **NOT** as an active topology selector.

     An ECMP-VID alone cannot select an active topology.

     ECMP-VID + MAC SA + other parameters are used for this
  4. And (carrying the properties of a B-VID) **NOT** as an end-to-end service identifier.

     I-SIDs are used for this.

BRIDGING EVOLUTION – IEEE802.1Qay

- An ESP-VID (subset of a B-VID) is used:
  1. As a control protocol identifier
  2. But **NOT** as VLAN connectivity identifier.
     An ESP-VID alone cannot identify the connectivity.
     TE-SID is used for this
  3. And **NOT** as an active topology selector.
     An ESP-VID alone cannot select an active topology.
     <ESP-DA, ESP-SA, ESPVID> is used for this
  4. And (carrying the properties of a B-VID) **NOT** as an end-to-end service identifier.
     I-SIDs are used for this.

CFM MA CONNECTIVITY IDENTIFIER

- Various connectivity identifiers corresponding to the classes of service instances the connectivity of which is to be monitored are provided for use by CFM
  - A list of VIDs for VLAN service instances
  - An I-SID for Backbone service instances
  - A TE-SID for TESIs
  - None of the above for untagged service instances
- Used by CFM PDUs in order to achieve “fate sharing” with the monitored service
- An “ECMP-SID” connectivity identifier is missing!

ECMP AND VLAN CONNECTIVITY

- Due to the SPBM VLAN symmetric properties (and its subset unicast/multicast congruency), CCMs on a SPBM VLAN can check the SPBM-VID connectivity using multicast frames which take the same path as any unicast in the same VLAN, LB and LT operates the way that one expects as MIPs are just processing CFM PDUs on the same connectivity path.
- Reference to ECMP-VID alone cannot identify connectivity. An ECMP-VID is only a control protocol identifier.

ECMP-SID

- Use a different identifier to introduce ECMP MAs.
- This could be represented by a new ECMP-SID parameter that would correspond to a tuple containing identifiers of the MA's end points and the SPBM-VID.
- For an ECMP path MA this would correspond to the potentially complete set of ECMP paths connecting the same set of two end points as contracted by the ECMP algorithm.
- For an ECMP VID MA this would correspond to the set trees connecting the same set of end points for a given SPBM-VID as constructed by ECMP.

CFM PDU ADDRESSING

- The destination addresses on CFM PDUs, monitoring a specific MA, are addresses that can reach the edges of that MA.
- In the case of MAs associated with services that do not permit flooding or learning these addresses can only be associated with the remote MEPs within the MA (as intermediate interfaces are invisible and are not reachable by forwarding functions supporting those services).
- For SPBM VLAN MAs such an address is the group MAC address for the SPBM default I-SID (Table 9-3, Figure 27-3)
- This address is not applicable for ECMP related services as the default I-SID is explicitly assigned to a symmetric ECT algorithm (28).
- For ECMP the address used must depend on the configuration and operation of ECMP as described in 27.17.2, 28.8.2 (unicast, multicast for source specific SPT, multicast for shared trees, head-end replication).

ECMP PATH MA

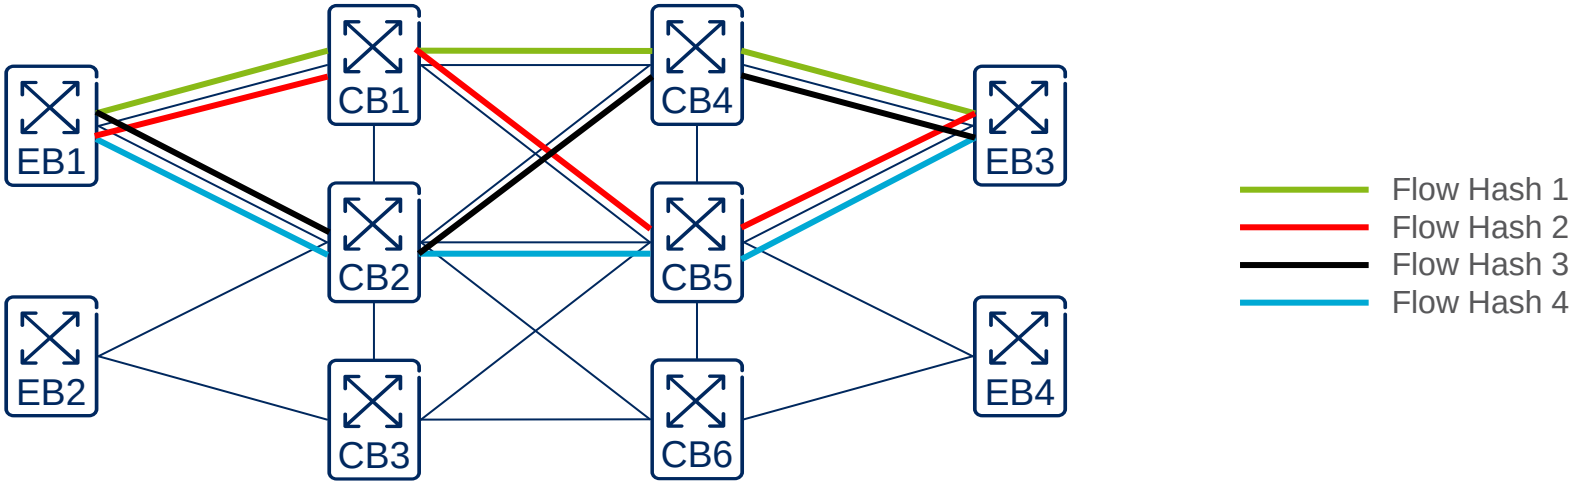

- ECMP Path MA identified by ECMP-SID: EB1, EB2, ECMP-VID [complete set of ECMP paths] -> This can be set by the operation of ECMP as soon as only EB1, EB3 are declaring the same I-SID
- CFM PDUs (DA, SA, VID) for EB1: (EB3, EB1, ECMP-VID)
- CFM PDUs DA for EB3: (EB1, EB3, ECMP-VID)
- CFM PDUs Flow Hash: cyclically set to all Flow Hash values

ECMP VID MA

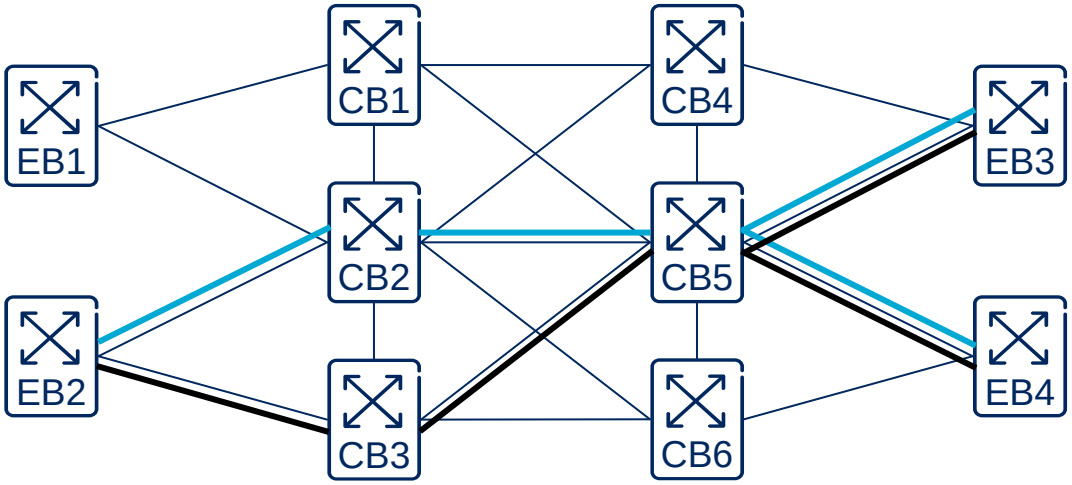

Shared Multicast 1
Shared Multicast 2

- ECMP VID MA identified by ECMP-SID: EB2, EB3, EB4, ECMP-VID [shared multicast] -> This can be set by the operation of ECMP as soon as only EB2, EB3, EB4 are declaring the same I-SID and the multicast for shared trees model is chosen (and the associated Bridge priorities are set [In the example depicted: in the light blue case CB2 is identified as the root, while in the black case CB3 is identified as the root]
- For every end-point in the multicast
  - CFM PDUs SA: Its own
  - CFM PDUs VID: ECMP-VID
  - CFM PDUs DA: cyclically set to all the Multicast Address identified by ECMP

SUMMARY

- ECMP MAs are different from PBB-TE and SPBM VLAN MAs as the monitored data service has different connectivity characteristics than that of TESIs or VLANs
- An ECMP path MA can be associated with a connectivity path connecting the associated endpoints or with a subset (not necessarily proper) of the equal cost paths connecting the same end points.
- The parameters on CFM PDUs for ECMP MAs depend on the configuration and operation of ECMP as described in 27.17.2, 28.8.2 (unicast, multicast for source specific SPT, multicast for shared trees, head-end replication).
- ECMP CFM PDUs are cyclically sent along all identified ECMP paths

ERICSSON