ECMP Interoperability in SPB

Ali Sajassi

November 11, 2010

IEEE 802.1 Plenary – Dallas, Tx

E-TAG Format

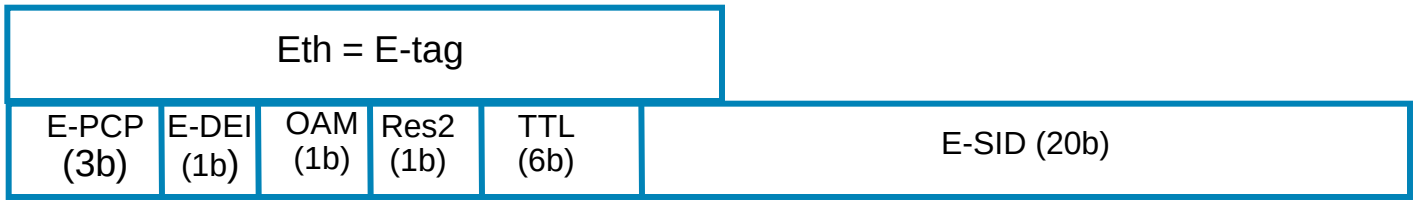

- E-Tag: ECMP Tag (TBA)
- E-PCP: same as I-tag
- E-DEI: same as I-tag
- OAM: 1-bit flag to indicate that this is an OAM frame
- Res2: 1-bit reserved
- TTL: 6-bit Time To Live
- E-SID: 20-bit version of I-SID

E-Tag

- Pros
  - Most efficient and compact encoding
  - 6-bit TTL accommodates both SP and DC applications
  - it is based on 802.1ah frame format that many vendors and providers are familiar with
- Cons
  - Reduces the Service ID field to 20 bits (but this doesn't create any inter-op issue as we will see)
  - Not backward compatible with existing BCB – requires BCBs to look at the TTL part of the E-tag (will be discussed later)
- NOTE: BCBs will only process TTL part of E-tag
  - E-SIDs are passed transparently just as before through BCBs

Inter-operability: Between Different Regions

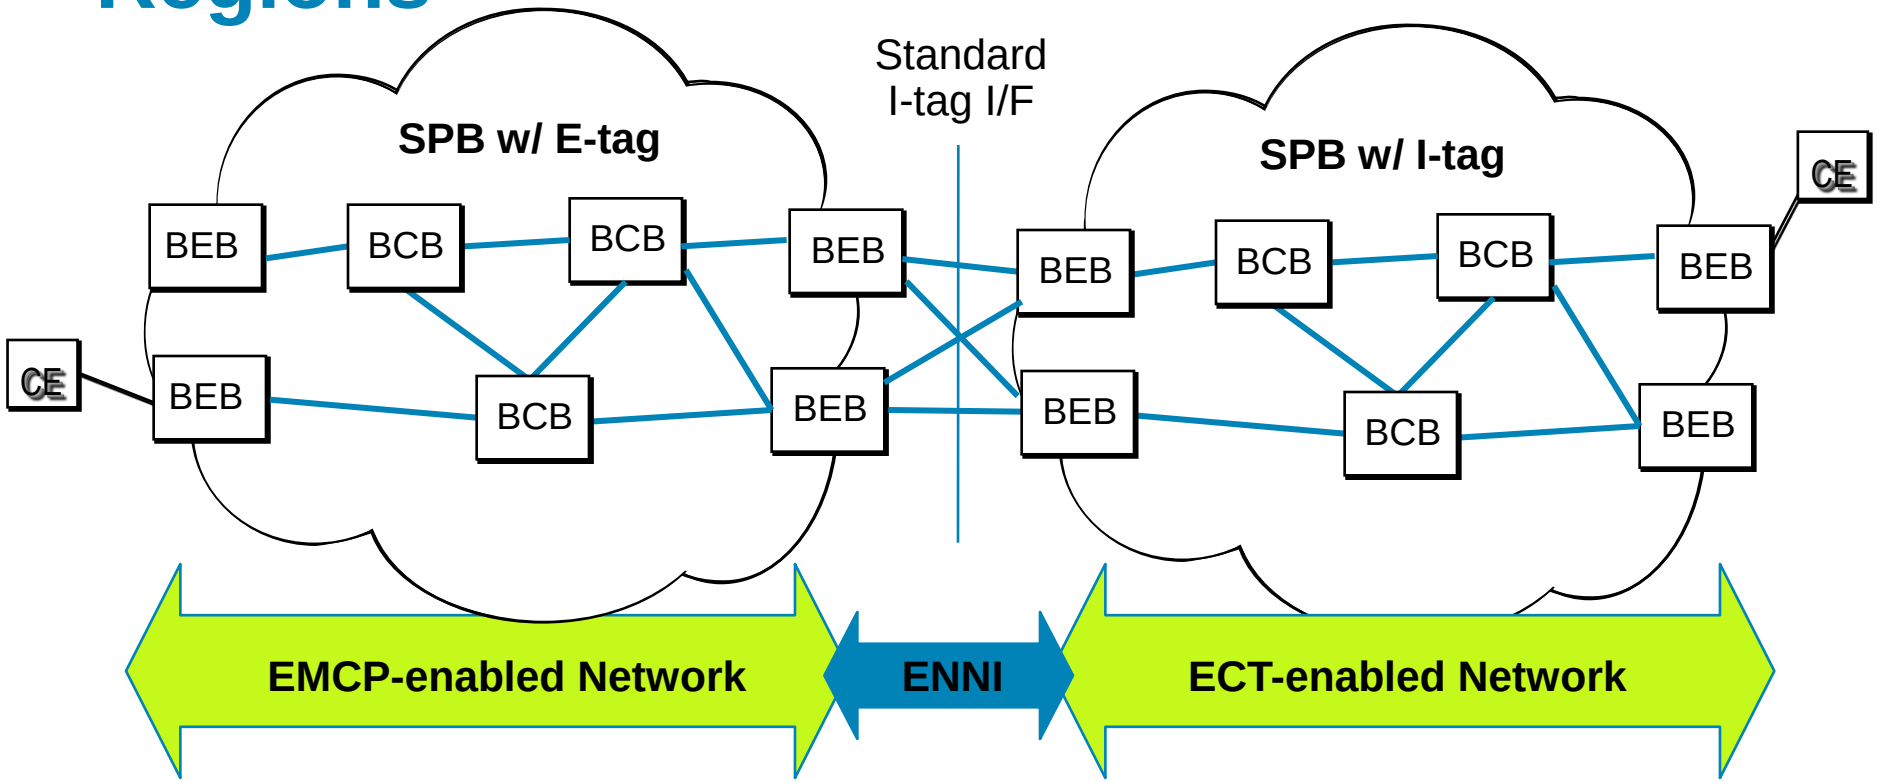

- BEB in the ECMP-enabled network will perform E-SID <-> I-SID mapping per existing 802.1ah functionality (per clause 6.11)
- BEB in the ECMP-enabled network will encode the derived I-SID into its corresponding I-tag and then send it to ECT-enabled network
- No changes is needed on the ECT-enabled network (both BEBs and BCBs)
- Number of I-SIDs supported over ENNI (using I-tag service interface) will be limited to 1 million instead of 16 millions (still much larger than any practical requirements) !!
- If needed to support 16 millions or more (upto 4 billions), then we can limit the scope of E-SID to B-VLAN

Inter-operability: Within one Region

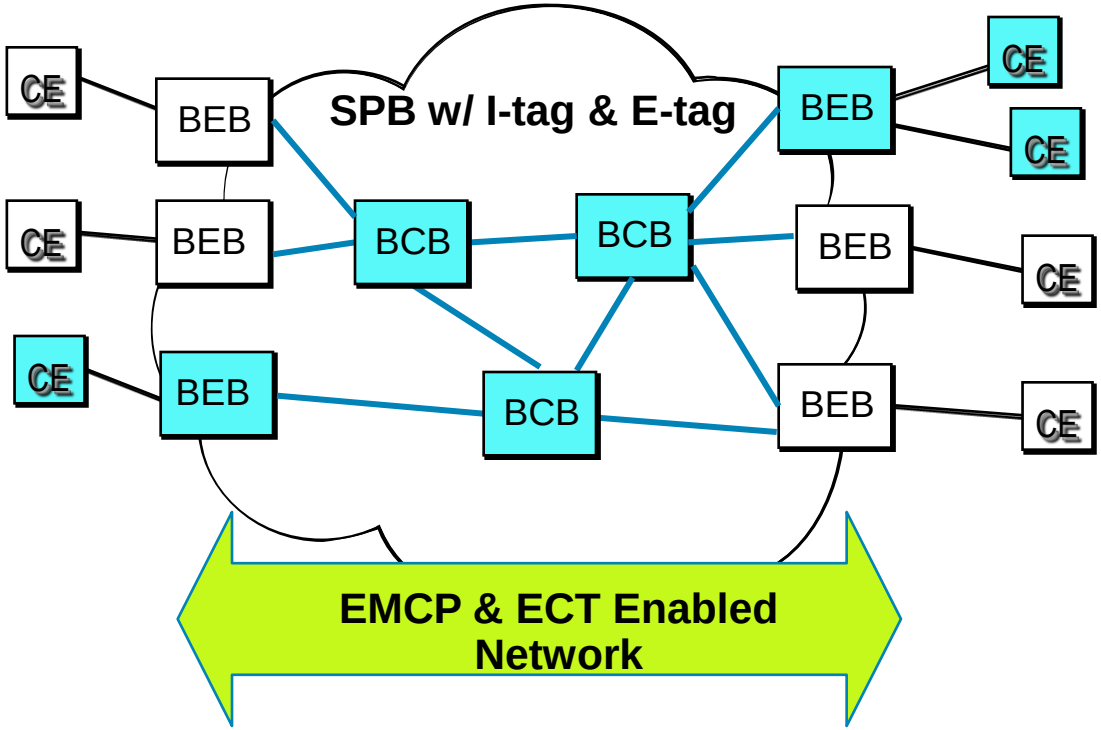

- Use designated B-VID(s) for ECMP just like
  - A set of B-VIDs for 802.1aq (one per ECT)
  - A set of B-VIDs for PBB-TE
  - A set of B-VIDs for PBB with MSTP
- To support ECMP
  - Some BCBs must support TTL
  - Only BEBs that are configured for E-SIDs, need to support TTL

Backward Compatibility

- Any per-hop ECMP (whether TTL is used or not) requires additional new processing anyway:
  - Hashing based on user data flow headers to determine egress interface or
  - using the pre-calculated hash-index to determine egress interface
- A network can be configured to simultaneously support ECMP and ECT modes
- In a single network, we cannot mixed ECMP service points with non-ECMP because it doesn't make sense
- In multiple networks where an ECMP service in one network needs to interoperate with non-ECMP service in another network, I-tag mapping capability of BEB can be used to ensure such interoperablity
- Multiple topology configuration can be used to support both ECMP and non-ECMP BCBs in the same network and ensure gradual migration