

802.1Qbp – ECMP Multicast Mode Selection January 2012

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

ECMP Behavior

- Default should spread both unicast and multicast traffic
- Random spreading should not require configuration
- Configuration knobs should be provided for additional control
 - Operator's choice whether or not to exercise this control
- Each SPB ECMP VLAN should be separately controllable
- Each BSI should be separately controllable

Unicast ECMP

- Random spreading hop-by-hop using hash of Flow ID, System ID
- Allows all possible equal cost paths to be used
- Does not guarantee particular distribution; but is deterministic

Multicast ECMP

- Each BSI endpoint must have a multicast mode
 - Source tree (unique group address per BSI endpoint)
 - Shared tree (one group address for all BSI endpoints)
 - Head end replication (no group address)
- Source tree calculation
 - Symmetry with unicast is no longer required
 - Randomize using per hop ECT tie-breaker or hash function
- Shared tree calculation
 - Single tree per root or multiple trees per root?
- How are BSI endpoints assigned to a shared tree
 - Random
 - Provisioned
 - Calculated

ECMP Multicast Attributes

- **Granularity of SPT selection?**
 - One (per region)
 - One per source node
 - N per source node
 - One per address
- **How many SPTs in selection set?**
 - One .1aq tie-breaker subset
 - N .1aq tie-breaker subsets
 - All SPTs
- **How many group addresses?**
 - One per I-SID endpoint
 - One per I-SID (requires shared tree)
- **Selection of SPT**
 - Automatic (requires standard hash)
 - Provisioned (may require ISIS-SPB extension)
- **Assignment of I-SID to SPT**
 - Automatic (requires standard hash)
 - Provisioned (may require ISIS-SPB extension)

Combinations in Draft 0.1

- **802.1aq** (first 16 ECT-ALGORITHMS)
 - One ECT tie-breaker, one SPT per source node
 - One address per I-SID endpoint
 - Fully automatic
- **Hash based source tree** (00-80-C2-21)
 - All SPTs possible, one SPT selected per source node (hash selects hop toward root)
 - One address per I-SID endpoint
 - Fully automatic
- **All ECT tie-breakers** (00-80-C2-22)
 - 16 ECT tie-breakers, 16 SPTs per source node
 - One address per I-SID endpoint
 - Fully automatic with provisioned override
- **Shared tree** (00-80-C2-22)
 - All SPTs possible, one selected per shared tree root node
 - One address per I-SID
 - Provisioned (or Automatic?)

Combination in “One Slide”

- **bp-ashwood-one-slide-1215-v3.pdf**
 - One ECT-ALGORITHM
 - 16 ECT tie-breakers
 - Source tree and shared tree supported
 - Source tree
 - 16 SPTs per source node possible
 - One SPT per source node by default (ECT tie-breaker 0)
 - One address per I-SID endpoint
 - Fully automatic with provisioned override
 - Shared tree
 - 16 shared tree roots (selected using 16 ECT tie-breakers)
 - One SPT per root, using same ECT tie-breaker
 - One address per I-SID
 - Automatic shared tree root selection with provisioned override
 - Provisioned I-SID assignment to shared tree (management automation possible)
- **Includes aspects of:**
 - **802.1aq** (but all ECT tie-breakers available in one VLAN)
 - **All ECT tie-breakers**
 - **Shared tree**

Multicast Mode Selection using ISIS-ADDR sub-TLV

| | | Octet | Length |
|------------------|----------------------------|--------------------|---------------|
| I-SID Tuple 1 | Type (3) | 1 | 1 |
| | Length | 2 | 1 |
| | B-MAC Address | 3-8 | 6 |
| | reserved | 9 | 4 bits |
| | Base VID | 9-10 | 12 bits |
| | T | 11 | 1 bit |
| I-SID Tuple n | R | 11 | 1 bit |
| | Multicast Selection | 11 | 6 bits |
| | I-SID | 12-14 | 3 |
| | ... | | |
| I-SID Tuple n | T | (4n+7) | 1 bit |
| | R | (4n+7) | 1 bit |
| | Multicast Selection | (4n+7) | 6 bits |
| | I-SID | (4n+8)- (4n+10) | 3 |

- Use reserved bits to select multicast mode option (if desired)
 - Ts (shared tree multicast source), ECT tie-breaker (4 bits)
- All zeros selects default behavior