

Control Plane Changes for 802.1Qbp

Ali Sajassi

May 17, 2011



Issue: I-SIDs to B-VID mapping

- Currently in clause 6.11, I-SIDs are groups into different B-VIDs bins
- In 802.1aq, for a given I-SID, the same B-VID is used for both unicast and mcast frames because both share the same ECT tree
- For ECMP operation, unicast and multicast frames need to use different trees and different load sharing algorithms
 - Unicast frames need to use ECMP algorithm
 - Multicast frames need to use ECT algorithm
- **Question: How should we do I-SID to B-VID mapping for ECMP ?**

Requirements:

- Use the same I-SID to B-VID mapping mechanism as currently specified in clause 6.11 – e.g., don't change the table (because it has hardware ramifications)
- Allow for different B-VIDs to represent different ECT algorithms as done currently with 802.1aq
- Allow for a single B-VID to represent different ECT algorithms in order to support untagged ECMP frames
- Allow for full ECMP operation of unicast frames regardless of whether multiple B-VIDs or a single B-VID is used
- Automatic distribution of I-SIDs among ECT algorithms is desirable
- Explicit distribution of I-SIDs among ECT algorithm (as done currently in 802.1aq) is mandatory
- Backward compatibility with existing 802.1aq specification is highly desirable
- Minimum changes to the spec. is highly desirable

Possible Options:

Options for I-SIDs to B-VIDs mapping:

- a) Use multiple B-VIDs to designate different ECTs as before for multicast frames and use the same set of B-VIDs for unicast ECMP frames
- b) Use a single B-VID to designate different ECT algorithms for multicast frames and use this B-VID for unicast ECMP frames
- c) Use multiple B-VIDs for multicast frames and use a single B-VID for unicast frames

Option A) Multiple B-VIDs for Both

- In this option different B-VID represent different ECT algorithm but all these B-VIDs are mapped to the same ECMP algorithm
- All B-VIDs are mapped to the same bridge domain (VLAN) – e.g., analogous to SVL
 - Same filtering DB is used for unicast (and multicast) addresses with different B-VIDs
 - SVL-like operation can be used because I-SIDs are unique across B-VIDs – e.g., same I-SID cannot be used under two different B-VIDs
- Pros
 - No modification to either control or data planes are needed
- Cons
 - We lose the option of not using B-tag for unicast ECMP because for a given I-SID, we need to use the same B-tag for unicast data as for multicast data

Option B)

- In this option a single B-VID is used to represent different ECT algorithms for multicast frames as well as the ECMP algorithm for unicast frames
- Pros
 - Use of B-tag can be optional for both unicast and multicast data
 - No additional changes in data plane is required
- Cons
 - It requires changes to the control plane – e.g., it requires decoupling of B-VID to ECT algorithm in SPBM I-SID sub TLV so that I-SIDs can be directly associated with ECT algorithm

Option C)

- In this option multiple B-VIDs are used to designate different ECT algorithms for multicast frames but a single B-VID is used to designate ECMP for unicast frames
- Pros
 - No changes to control plane is required
 - Use of B-tag can be optional for unicast frames
- Cons
 - B-tag must always be used for multicast frames
 - It requires changes to data-plane to associate two B-VIDs for the same I-SID (one for unicast and the other for multicast)

Changes Needed for Option (b)

- Option (b) requires no changes to data-plane and minor changes to control plane
 - SPBM ISID-ADDR TLV should be modified to replace base VID with ECT Algorithm ID so that I-ISIDs are directly associated with the algorithm ID
- It is highly desirable to ensure backward compatibility and to avoid introduction of any new TLV/sub-TLV or modification of existing TLV/sub-TLV for ECMP operation
- In order to do this:
 - We will use the same Base VID field in SPBM ISID-ADDR TLV for ECT Algorithm ID
 - We use one bit of 4-bit reserved field to indicate whether the next 12 bits is a “Base VID” or is a “ECT Algorithm ID” (0= base vid, 1= alg ID). The usage of the reserved field in this way, guarantees backward compatibility
 - In order to fit the ECT algorithm ID in a 12-bit field, we use the 8-bit index part of the “ECT algorithm ID” and assume that OUI is the standard one

Association of I-SIDs to a ECT Algorithm (indirectly via Base VID)

		Octet	Length
		1	1
		2	1
		3-8	6
		9	4 bits
		9-10	12 bits
I-SID Tuple 1	T	11	1 bit
	R	11	1 bit
	reserved	11	6 bits
	I-SID	12-14	3
...			
I-SID Tuple n	T	(4n+7)	1 bit
	R	(4n+7)	1 bit
	reserved	(4n+7)	6 bits
	I-SID	(4n+8)- (4n+10)	3

1-bit flag
 0= Base VID
 1= Algorithm ID

V/A (1bit)

Base VID /Alg ID

Figure 28-11—SPBM Service Identifier and Unicast Address sub-TLV

ECT Algorithm sub-TLV

No Changes is required to this TLV

	Octet	Length
Type (TBD)	1	1
Length	2	1
ECT ALGORITHM	3-6	4
ECT Information	7-(Length+2)	variable

- a) Type (8-bit) Value TBD
- b) Length (8-bits)
Total number of bytes contained in the value field.
- c) ECT-ALGORITHM (4-bytes)
ECT-ALGORITHM is advertised when the bridge supports a given ECT-ALGORITHM (by OUI/Index) on a given VID.
- d) ECT Information (variable)
ECT-ALGORITHM Information of variable length.

Mapping between ECT Algorithm & Base VID

No Changes is required to this TLV

		Octet	Length
ECT-VID Tuple 1	Type (TBD)	1	1
	Length (6n)	2	1
	ECT Algorithm	3-6	4
	Base VID	7-8	12 bits
	U	8	1 bit
	M	8	1 bit
	reserved	8	2 bits
...			
ECT-VID Tuple n	ECT Algorithm	(6n-3)-6n	4
	Base VID	(6n+1)- (6n+2)	12 bits
	U	6n+2	1 bit
	M	6n+2	1 bit
	reserved	6n+2	2 bits

Figure 28-5—SPB Base VLAN-Identifiers sub-TLV

Example: 802.1aq & .1Qbp option (a)

ECTs	Algorithm	B-VID
1	Tie—break 1	11
2	Tie—break 2	12
3	Tie—break 3	13
4	Tie—break 4	14
5	Tie—break 5	15

In Control Plane

- No changes to control plane – operate as before
- Use different tie-breaking algorithm for different ECTs
- Assign a base VID for each of these algorithms
 - SPB base VID sub-TLV
- Load share I-SIDs among different ECTs and associated I-SID to B-VID
 - use SPBM I-SID sub-TLV as is w/o any changes
 - I-SID to B-VID association is derived directly by getting I-SID -> B-VID association from SPBM I-SID sub-TLV

In Data Plane

- Program Filtering DB for each mcast group address based on the above I-SID to B-VID association

Example: .1Qbp option (b)

ECTs	Algorithm	B-VID
1	Tie—break 1	11
2	Tie—break 2	11
3	Tie—break 3	11
4	Tie—break 4	11
5	Tie—break 5	11

In Control Plane

- Minor change in control plane: set V/A flag in SPBM I-SID sub-TLV to 1 & copy the algorithm index in the base VID filed
- Use different tie-breaking algorithm for different ECTs
- Assign a **single base VID** for all of these ECTs algorithms in SPB base VID sub-TLV
- Load share I-SIDs among different ECTs and associated I-SID to B-VID
 - I-SID to B-VID association is derived indirectly by first getting I-SID -> Alg-ID association from SPBM I-SID sub-TLV and then getting Alg-ID to B-VID association from SPB Base VID sub-TLV

In Data Plane

- Program Filtering DB for each mcast group address based on the above I-SID to B-VID association

Recommendation

- Use option (b) whenever it is required to use a single B-VID for all ECT algorithms designation of multicast data
- Use option (a) whenever it is required to use multiple B-VIDs for ECT algorithms designation of multicast data (one B-VID per ECT algorithm)
- The choice between option (a) and option (b) needs to be done network wide for 802.1Qbp operation
 - They should be mutually exclusive for a given .1Qbp network
- Multi-topology still uses different B-VIDs – e.g., in case of option (b), different B-VIDs can designate different MT instances

