

IEEE 802.1Qbf Three Issues for
Discussion
March 2010 Interim Meeting
Orlando, Florida, USA

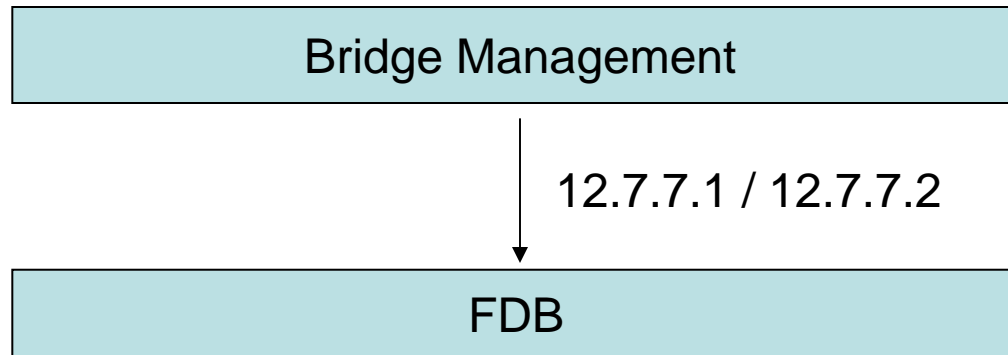
Bob Sultan (bsultan@huawei.com)

Ben Mack-Crane

Three Qbf Issues


1. How to create/delete Static FDB entries when IPS is deployed?
2. An IPG should be associated with a list of 2-tuples $\langle \text{ESP_DA}, \text{ESP_VID} \rangle$ rather than a list of TESIs.
3. Description of the Segment MA (just to make sure everyone is in agreement).

Issue 1: Create/delete Static FDB Entries



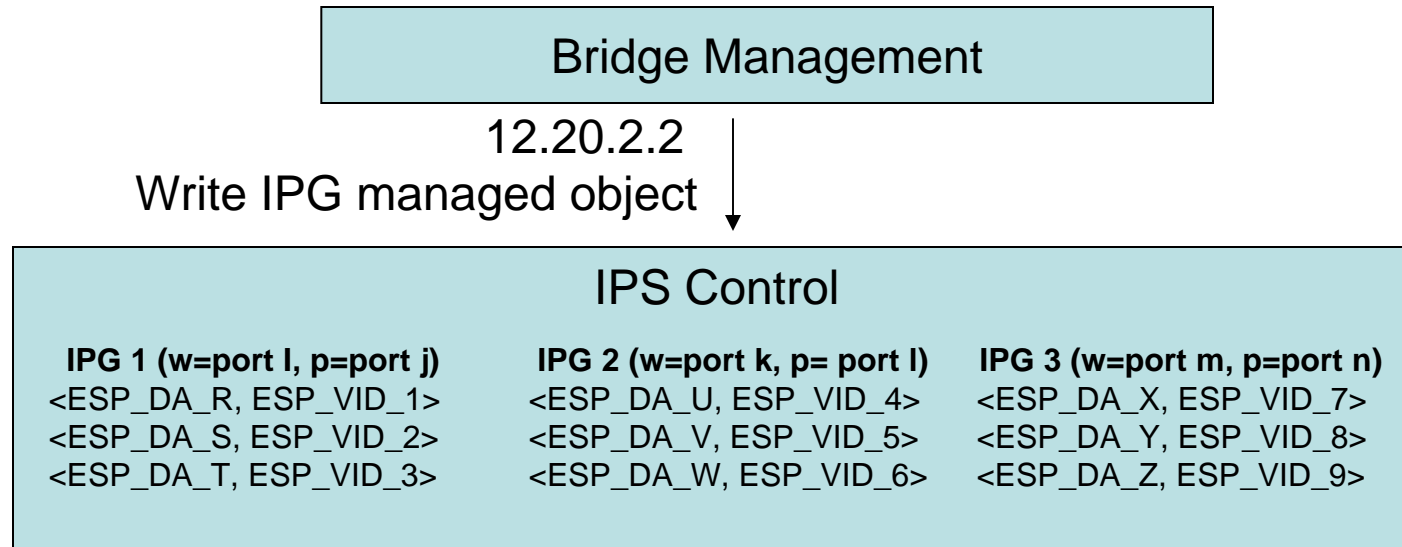
- An ESP is associated with a Static FDB entry, identified by the 2-tuple $\langle \text{ESP_DA}, \text{ESP_VID} \rangle$, in each FDB through which it passes;
- 12.7.7.1 Create Filtering Entry describes the creation of such an FDB entry by Bridge Management;
- 12.7.7.2 Delete Filtering Entry describes the deletion of such an FDB entry by Bridge Management;

Possible Race When IPS Introduced



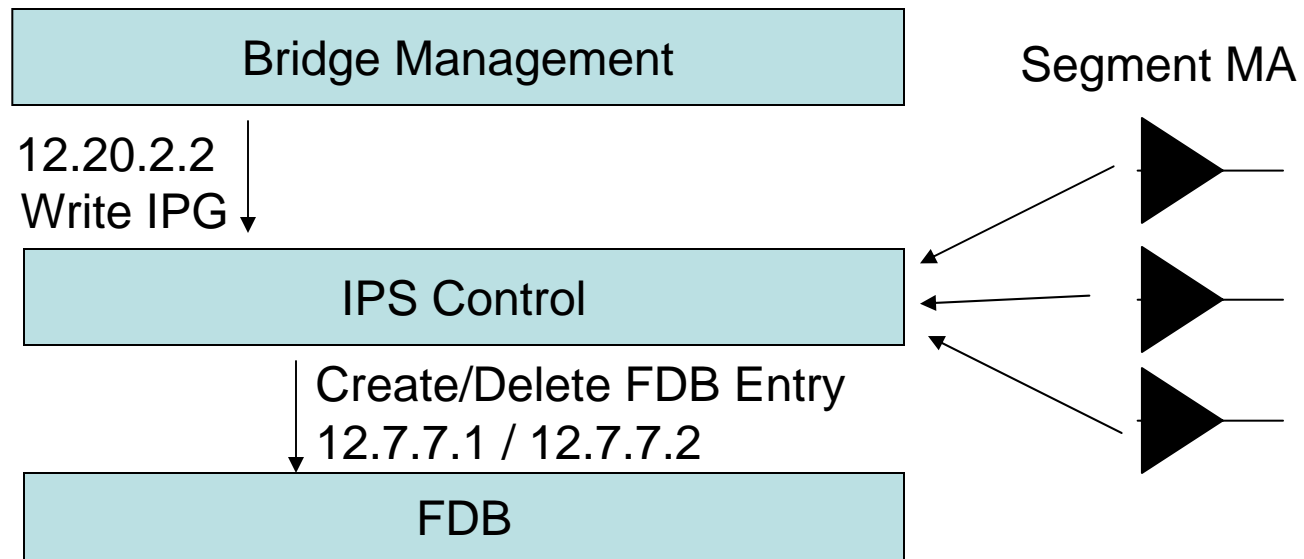
- IPS control must be able to change the value of the outbound Port field of an FDB entry depending on the state of the IPS State Machines (Working or Protect);
- A race condition can occur if Bridge Management deletes FDB entry X at the same time that IPS Control updates the outbound Port value of FDB entry X;

Avoiding Race Condition



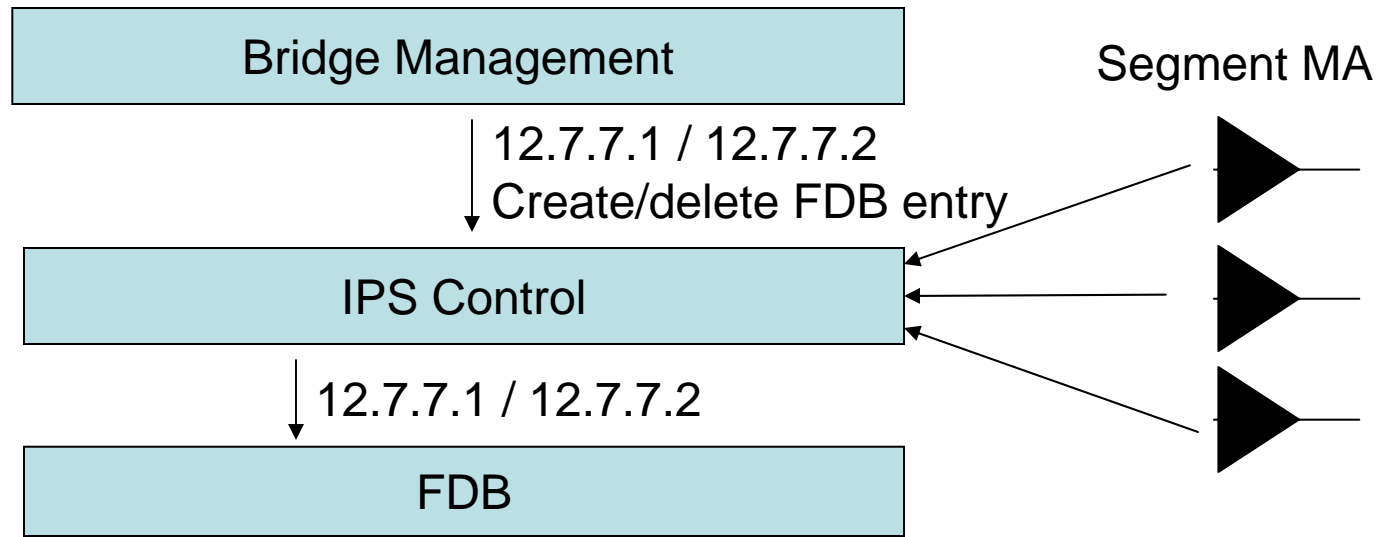
- Bridge Management communicates to IPS Control (via 12.20.2.2 Write IPG) the list of 2-tuples identifying FDB entries associated with each IPG for which the Bridge serves as the IPG Ingress;
 - i.e., these are exactly the FDB entries whose outbound Port value is maintained by IPS control (using the IPS State Machines);

Avoiding Race Condition (con'd)



- If a 2-tuple is added to the list, IPS Control issues Create FDB Entry to the FDB with an outbound Port value equal to the Working SEID or the Protection SEID, depending on the state of the IPS State Machines;
- If a 2-tuple is deleted from the list, IPS Control issues Delete FDB Entry to the FDB;
- If a 2-tuple is currently on the list, and the state of the associated IPS State Machine changes; IPS control issues Create FDB Entry to change the outbound Port value;

Avoiding Race Condition (con'd)



- On receiving create/delete FDB Entry from Bridge Management, IPS control determines whether the specified FDB entry is on an IPG 2-tuple list;
 - i.e., was previously added to an IPG by 'Write IPG';
- If so, IPS Control fails the requested create/delete FDB entry (IPS Control 'owns' this FDB entry);
- If not, IPS Control passes the request on to the FDB without modification (Bridge Management 'owns' the FDB entry); ⁷


Issue 1 conclusion

- The method described is not *exactly* the method described in D0.2;
- The editor will correct this in the next draft if there is agreement on this solution;

Issue 2: List associated with an IPG

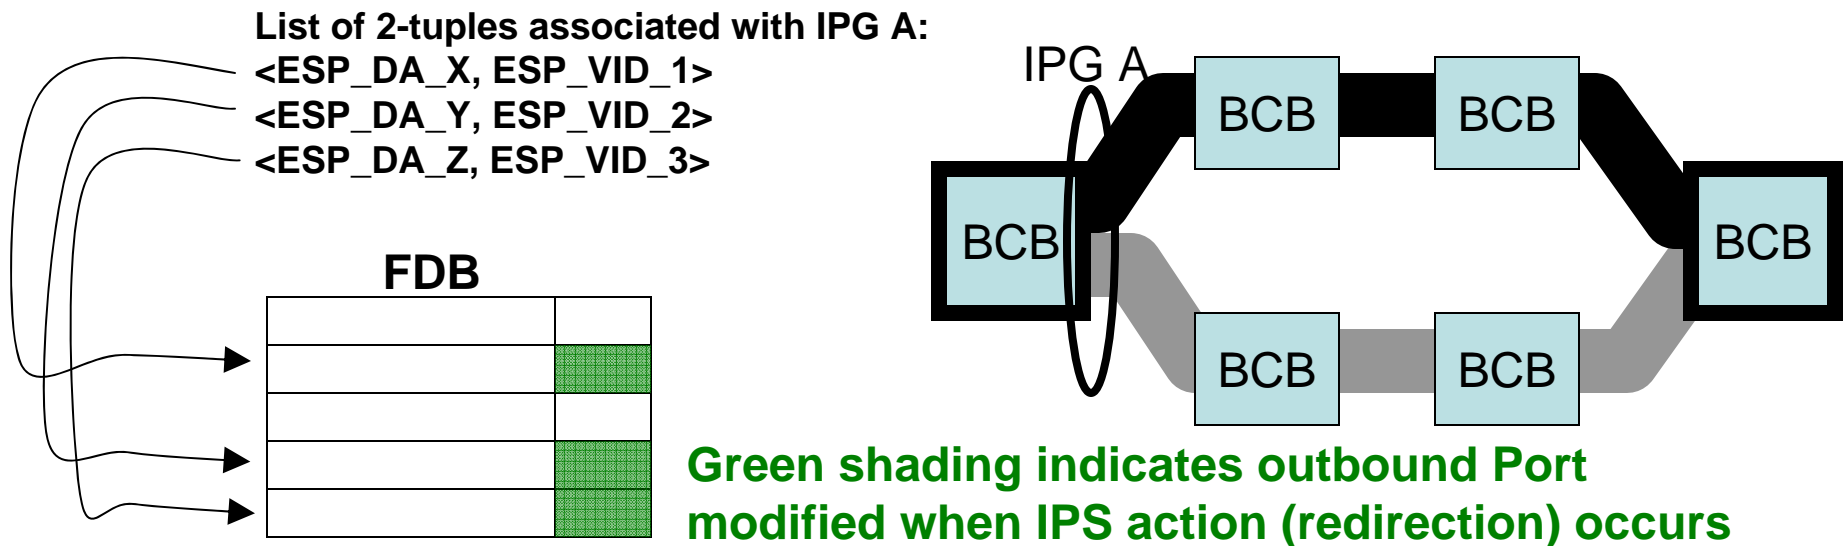
- D0.2 describes that an IPG is associated with a list of *TESIs* for which redirection is performed when an IPS action occurs;
- When an IPS action occurs, at an IPG Endpoint, one steps through the list of *TESIs*;
 - Determines which ESP in the pair is pointing towards the IPG;
 - Finds the FDB entry having a 2-tuple $\langle \text{ESP_DA}, \text{ESP_VID} \rangle$ matching the 3-tuple $\langle \text{ESP_DA}, \text{ESP_SA}, \text{ESP_VID} \rangle$ identifying the ESP;
 - Updates the 'outbound Port' field FDB with the SEID of the Active Segment for the IPG (determined by the IPS State machines for the IPG);
 - It is possible that the same FDB entry will be updated (with the same outbound Port value) multiple times.

A Picture of what was just described



- The figure shows one IPG (Black/Gray);
- The BCBs that are the IPG Endpoints have a thick outline;
- Three TESIs (red, blue, green) are associated with the IPG at each IPG Endpoint;
- ***But this is very odd because there is no other case in which a TESI is provisioned on a BCB;***
- ***Nor does a BCB have any knowledge of TESIs***

The Fix



- An IPG Endpoint need have no knowledge of TESIs or of ESPs;
- The IPG Endpoint need only have the list of 2-tuples identifying the FDB entries whose outbound Port fields are updated when an IPS action occurs;
- This is *entirely* consistent with current .1Q architecture ***which does not require any provisioning of TESIs or ESPs in a BCB;***
- One simply provisions the appropriate Static FDB Entries

Conclusion

- The requirement to provision TESIs in a BCB is inconsistent with existing PBB-TE architecture;
- Conceptually, it is exactly a list of 2-tuples <ESP-DA, ESP-VID> that we want to associate with an IPG as this is what allows identification of FDB entries whose outbound Port entries are to be updated when an IPS action occurs;
- In the absence of evidence to the contrary, the next draft will be updated to reflect this;

Issue 3: The Segment MA

- Three types of MA are currently defined:
 1. VLAN-based (identified by VID);
 2. Backbone service instance based (identified by I-SID);
 3. PBB-TE MA – associated with TESI (identified by TE-SID);
- We add fourth type of MA, associated with a *Segment* rather than a *Service Instance*;
- Organize four types of MA like this:
 1. VLAN-based (identified by VID);
 2. Backbone service instance based (identified by I-SID);
 3. PBB-TE MA
 - a) TESI MA – associated with TESI (identified by TE-SID);
 - b) Segment MA – associated with Segment (identified by TE-SID; i.e., pair of 3-tuples);

Properties of the Segment MA

- Two types of PBB-TE MA differ in that
 - TESI MA is instantiated on CBP;
 - Segment MA is instantiated on PNP
- Although Segment MA is identified by a TE-SID, the Segment MA is not associated with a TESI;
 - The TE-SID is simply an identifier;
- Although the TE-SID comprises a pair of 3-tuples; the 3-tuple is not associated with an ESP;
 - The 3-tuple is simply an identifier;

Why organize the MA types like this?

- There are many references to PBB-TE MA in clauses 19 – 22;
- Most of these references need not be changed;
- In cases where behavior of TESI MA and Segment MA are different, text is changed to explicitly specify TESI MA or Segment MA and the associated behavior;
- D0.2 currently uses this approach (so this discussion is simply to ensure that everyone is in sync)