8. Principles of bridge operation

This proposed text follows detailed review of P802.1ay D4.0. I found that draft rather loose in the way it uses or bends the concepts of active topology and Port State, and not accurate in detail when it uses those as part of its specification. In particular D4.0's changes to 8.6.1 claim "The Forwarding Process allocates each received frame to an active topology". This is simply not possible for ESPs (at least not as part of 8.6.1) since the active topology for ESPs is really implemented in the FDB - so it is not known at stage 8.6.1. So far as 8.6.1 is concerned *forwarding* is FALSE and *learning* TRUE for all frames with an ESP-VID. This shows that D4.0 8.4 is incorrect when it says "forwarding is managed by an external agent" in the context of ESPs and Port State. It is true that the forwarding done by the Bridges as a whole is managed by an external agent, but the Port State/VID related part of that is not. The description in clause 8 has to work as an overall model, not just as a set of observations.

P802.1ay has to be properly integrated into clause 8 and I make no claim that the changes here cover all the issues. For example it is necessary, for the Filtering Database to work correctly, that PBB-TE specify the mapping of VIDs to FIDs. Anyone who thinks that this is unneccessary because PBB-TE doesn't do learning needs to read clause 8.8 again.

Replace clause 8.4 with the following:

<These changes are marked up using 802.1Q-2006 edition 0.1 as a base. Since paragraphs and sentences have been reordered, even though changes to the text within those paragraphs and sentences are shown, the editing instruction is 'replace' rather than 'change'.>

8.4 ~~Port states and the active topology~~Active topologies, learning, and forwarding

~~Each Bridge Port has an operational Port State that governs whether or not it forwards frames classified as belonging to a given VLAN, and whether or not it learns from their source addresses.~~

~~The~~ An *active topology* ~~of a Bridged Local Area Network at any time is the set of communication paths formed~~ is the connectivity provided, for frames with a specified set of VID, destination address, and source address values, ~~a set of communication paths formed~~ by interconnecting the LANs and ~~B~~bridges in a bridged network ~~by the~~ with *forwarding* Bridge Ports.

~~The function of the distributed Spanning Tree algorithm and the Rapid Spanning Tree Protocol (RSTP, IEEE Std 802.1D Clause 17) used by SST Bridges to execute that algorithm, is to construct an active topology that is simply connected relative to communication between any pair of end stations, irrespective of the VLAN classification of frames used. The Multiple Spanning Tree Protocol (MSTP, Clause 13) used by MST Bridges constructs multiple active topologies,~~ The distributed spanning tree algorithms and protocols, i.e. the Rapid Spanning Tree Protocol (RSTP, IEEE Std 802.1D Clause 17), and the Multiple Spanning Tree Protocol (MSTP, Clause 13), construct one or more active topologies, each simply and fully connected for frames belonging to any given VLAN. The *forwarding* and *learning* performed by each Bridge Port for each spanning tree is dynamically managed by RSTP, or MSTP to prevent temporary loops and reduce excessive traffic in the network while minimizing denial of service following any change in the *physical topology* of the network.

Provider Backbone Bridge Traffic Engineering enables construction of active topologies by the external agent that is responsible for setting up provisioned Ethernet Switched Paths (ESPs).

Any port that is not enabled, i.e., has MAC_Operational (6.6.2) False or has been excluded from the active topology by management setting of the Administrative Bridge Port State to Disabled, ~~or has been dynamically excluded from forwarding and learning from MAC frames, is assigned the Port State Discarding~~ has both *forwarding* and *learning* disabled for all spanning trees and ESPs.

If the Bridge Port is enabled, PBB-TE disables *learning* and enables *forwarding* for all frames allocated to each ESP-VID. An external agent manages the Filtering Database to control the forwarding of frames with particular values of ESP-VID and destination MAC address.

The term *Port State* summarizes per tree combinations of *forwarding* and *learning*, and any additional per tree variables used by a given spanning tree protocol to enforce the active topologies it has calculated, and is used by RSTP and MSTP as follows. Any port that has learning and forwarding disabled is assigned the Port State *Discarding* ~~Any~~, a Port that has learning enabled but forwarding disabled ~~for frames allocated to a given spanning tree~~ has the Port State *Learning* ~~for that tree~~, and a Port that both learns and forwards frames ~~if~~ has the Port State ~~is~~ *Forwarding*. However the RSTP and MSTP state machines (IEEE Std 802.1D-2004, Clause 17 and Clause 13 of this standard) do not control the Port State directly but use independent *forwarding* and *learning* variables for each tree.

RSTP constructs a single spanning tree, the Common Spanning Tree (CST), and maintains a single Port State for each Bridge Port. SST Bridges allocate all frames to that single spanning tree irrespective of their VLAN classification or source and destination MAC addresses.

MSTP constructs multiple spanning trees, the Common and Internal Spanning Tree (CIST) and additional Multiple Spanning Tree Instances (MSTIs), and maintains a Port State for each spanning tree for each Port. An MST Bridge allocates all frames classified as belonging to a given VLAN to the CIST or to one of the MSTIs using the MST Configuration Table. An MISTID of 0 identifies the CIST. A reserved MSTID value (TE-MSTID) is used to identify VIDs for use by PBB-TE with ESPs. A single VID is used to identify frames assigned to any given VLAN that is supported by the CIST or an MSTI. That VID is used by the Forwarding Process (8.6) to identify the spanning tree for the relayed frame, and thus identifies the applicable Port State.

Figure 8-5 illustrates the operation of the Spanning Tree Protocol Entity, which operates the Spanning Tree algorithm and its related protocols, and its modification of Port state information as part of determining the active topology of the network.

Figure 8-3 illustrates the Forwarding Process's use of the Port State: first, for a Port receiving a frame, to determine whether the received frame is to be relayed through any other Ports; and second, for another Port in order to determine whether the relayed frame is to be forwarded through that particular Port.

Figure 8-4 illustrates the use of the Port state information for a Port receiving a frame, in order to determine whether the station location information is to be incorporated in the Filtering Database.

Replace clause 8.6.1 with the following:

8.6.1 Active topology enforcement

~~The Forwarding Process allocates each received frame to a spanning tree.~~ To prevent data loops and unwanted learning of source MAC addresses, the Forwarding Process determines the values (True, or False) of the *learning* and *forwarding* controls (8.4) appropriate to each received frame and Bridge Port. If *learning* is True for the receiving Port and ingress filtering (8.6.2) would not cause the received frame to be discarded, ~~If the reception Port State for that spanning tree is Forwarding or Learning,~~ the frame's source address and VID are submitted to the Learning Process. If *forwarding* is True for the receiving Port, ~~If the reception Port State is Forwarding,~~ each Bridge Port, other than that receiving Port, with *forwarding* True ~~with a Port State of Forwarding for that tree~~ is identified as a potential transmission Port.

An SST Bridge ~~allocates all frames to a single spanning tree~~ supports a single active topology, the Common Spanning Tree (CST). For each Bridge Port, RSTP determines a single value for *forwarding* and a single value for *learning* (8.4) for all frames.

Bridges with MST, PBB-TE, or SPB capabilities, use the VID of the received frame to determine *forwarding* and *learning* for each Bridge Port for that frame as follows.

If the bridge supports PBB-TE, and the VID is an ESP-VID, *forwarding* is True and *learning* False. Control over the active topology of each ESP is provided by configuration of static filtering entries in the Filtering Database (8.8.1, 25.10.1).

If the bridge uses MSTP to configure the CIST, and the VID identifies a VLAN assigned to the CIST, *forwarding* and *learning* are as determined by MSTP (Clause 13) for the CIST.

If the bridge uses MSTP to configure MSTIs, and the VID identifies a VLAN assigned to an MSTI calculated by MSTP, *forwarding* and *learning* are as determined by MSTP for that MSTI.

All bridges other than SST Bridges implement the MST Configuration Table (8.9.1). The use of VIDs to determine learning and forwarding, as required by this clause shall be consistent with that table as follows. All VIDs allocated by the MST Configuration Table to the CIST-MSTID (0x000) are assigned to the CIST, all VIDs allocated to the TE-MSTID (0xFFE) are ESP-VIDs, and all VIDs allocated to other values of MSTID are assigned to the MSTI identified by that MSTID.

~~An MST Bridge allocates all frames with a given VID to the CIST or to a Multiple Spanning Tree Instance (MSTI). The allocation can be controlled by configuration of the MST Configuration Table (8.9.1) maintained by the Forwarding Process, subject to constraints (if any) imposed by the allocation of VIDs to FIDs (8.8.7). VIDs allocated to different spanning trees shall also be allocated to different FIDs. VIDs allocated to a given spanning tree may share the same FID.~~

Change the heading of clauses 8.6.2 as follows:

8.6.2 Ingress filtering

Change bullet (c) of clause 8.6.7 as follows.

c) If the ~~associated~~ transmission Port is no longer *forwarding* for that frame (8.4, 8.6.1) ~~leaves the Forwarding state~~.

8.7 The Learning Process

Change the initial paragraph of clause 8.7, creating two paragraphs as follows.

The Learning Process receives the source MAC Addresses and VIDs of received frames from the Forwarding Process, subject to active topology enforcement (8.6.1) and ~~following~~ the application of ~~the~~ ingress ~~rules~~ filtering (8.6.2). The Learning Process is not invoked (see 8.6.1) for frames whose VID is an ESP-VID or identifies a B-VLAN supported by SPBB.

When invoked, the Learning Process ~~It~~ shall create or update a Dynamic Filtering Entry (8.8.3) that specifies the reception Port for the frame's source address and VID, if and only if the source address is an Individual Address, i.e., is not a Group Address, the resulting number of entries would not exceed the capacity of the Filtering Database, and the filtering utility criteria for the receiving Bridge Port are met, as specified below.

8.8 The Filtering Database

Change the initial text of clause 8.8 as follows.

The Filtering Database supports frame filtering (8.6.3) queries by the Forwarding Process ~~as~~ to determine whether received frames ~~received by the Forwarding Process~~, with given values of destination MAC Address ~~parameter~~ and VID, are to be forwarded through a given potential transmission Port (8.6.2, 8.6.3, Delete the clause heading "8.8.7.1 Fixed and dynamic VID to FID allocations", adding the text of that clause after the text of 8.8.7, with the following changes:

8.8.8 Querying the Filtering Database

Retain the existing Tables 8-5 and 8-6 referenced by clause 8.8.8 unchanged, but change the text of the clause as follows:

~~If a frame is classified into a VLAN containing a given outbound Port in its member set (8.8.9), forwarding or filtering through that Port is determined by the control elements of filtering entries for the frame's destination MAC Address and for VLANs with the same VID or Filtering Identifier (FID, 8.8.7) as the frame's VLAN.~~

If the VID assigned to a relayed frame identifies a VLAN with a given outbound Port in its member set (8.8.9), the Filtering Database entries that are applicable to the frame's destination MAC Address and VID determine whether the frame is filtered or forwarded through that Port. More than one entry can apply to a given frame. This clause (8.8.8) specifies the effect of combining applicable entries, and of the absence of certain types of entries (not all possible entries are necessarily present). For an overview of the different types of entries see the introductory text of clause 8.8.

~~Each entry in the Filtering Database for a MAC Address comprises~~

~~a) A MAC Address specification;~~
~~b) A VID or, in the case of Dynamic Filtering Entries, an FID;~~
~~c) A Port Map, with a control element for each outbound Port.~~

~~For Dynamic Filtering Entries, the FID that corresponds to a given VID is determined as specified in 8.8.7.~~

~~For a given VID, a given individual MAC Address specification can be included in the Filtering Database in a Static Filtering Entry, a Dynamic Filtering Entry, both or neither. Table 8-5 combines Static Filtering Entry and Dynamic Filtering Entry information for an individual MAC Address to specify forwarding, or filtering, of a frame with that destination MAC Address and VID through an outbound Port.~~

~~NOTE 1—The use of FID in this table for Static Filtering Entries, and the text in parentheses in the headings, reflects the fact that, where more than one VID maps to a given FID, there may be more than one Static Filtering Entry that affects the forwarding decision for a given individual MAC Address. The effect of all Static Filtering Entries for that address, and for VIDs that correspond to that FID, is combined, such that, for a given outbound Port:~~
~~IF <any static entry for any VIDs that map to that FID specifies Forwarding> THEN <result = Forwarding>~~
~~ELSE IF <any static entry for any VIDs that map to that FID specifies Filtering> THEN <result = Filtering>~~
~~ELSE <result = Use Dynamic Filtering Information>~~

The Filtering Database entries applicable to a frame whose destination MAC Address is a specific individual MAC Address are the Dynamic Filtering Entry (if present) with that MAC Address and the FID to which the frame's VID is allocated (8.8.7), and all Static Filtering Entries (if any) with that MAC Address and a VID that is also allocated to that FID. An entry with a wildcard VID applies only if there is no applicable Static Filtering Entry for a specific VID. Table 8-5specifies the result of combining this informationfor an individual MAC address and a given outbound Port, and can be summarized as follows:

— IF any static entry for a VID allocated to the FID specifies Forward THEN Forward;
— ELSE IF any static entry for a VID allocated to the FID specifies Filter THEN Filter;

- ELSE IF a static entry for the wildcard VID specifies Forward THEN Forward;
- ELSE IF a static entry for the wildcard VID specifies Filter THEN Filter;
- ELSE IF dyamic (learnt filtering information) entry for the FID specifies Filter THEN Filter;
- ELSE Forward.

~~Table 8-6 specifies the result, Registered or Not Registered, of combining a Static Filtering Entry and a Group Registration Entry for the "All Group Addresses" address specification, and for the "All Unregistered Group Addresses" address specification for an outbound Port. Table 8-7 combines Static Filtering Entry and Group Registration Entry information for a specific group MAC Address with the Table 8-6 results for All Group Addresses and All Unregistered Group Addresses to specify forwarding, or filtering, of a frame with that destination group MAC Address through an outbound Port.~~

The Filtering Database entries applicable (if present) to a frame whose destination MAC Address is a specific group MAC Address are the Static Filtering Entries and Group Registration Entries (for the frame's VID and for the wildcard VID) whose address specification is that group MAC address or "All Group Addresses" or "All Unregistered Group Addresses". A Static Filtering Entry for a wildcard VID applies only if there is no applicable Static Filtering Entry for a specific VID, Group Registration Entries do not use wildcard VIDs. Table 8-6 specifies the results, Registered or Not Registered for a given outbound Port, of combining the entries for the "All Group Addresses" and for combining the entries for "All Unregistered Group Addresses". Table 8-7 combines these results with the entries for the specific group MAC Address. The result of Table 8-6 for "All Group Addresses" can be summarized as follows:

- IF a static entry for "All Group Addresses" and the frame's Base VID specifies Forward (Registration Fixed) THEN "All Group Addresses" is Registered;
- ELSE IF a static entry for "All Group Addresses" and the frame's Base VID specifies Filter (Registration Forbidden) THEN "All Group Addresses" is Not Registered;
- ELSE IF a static entry for "All Group Addresses" and the wildcard VID specifies Forward (Registered) THEN "All Group Addresses" is Registered;
- ELSE IF a static entry for "All Group Addresses" and the wildcard VID specifies Filter (Registration Forbidden) THEN "All Group Addresses" is Not Registered;
- ELSE IF a dynamic (Group Registration) entry for "All Group Addresses" and the frame's Base VID specifies Forward (Registered) THEN "All Group Addresses" is Registered;
- ELSE "All Group Addresses" is NOT Registered

NOTE—The result for "All Unregistered Group Addresses" is given by substituting that address specification for "All Group Address" throughout the summary.

The result of Table 8-7 can be summarized as follows:

- IF a static entry for the specific group address and the frame's Base VID specifies Forward THEN Forward;
- ELSE IF a static entry for the specific group address and the frame's Base VID specifies Filter THEN Filter;
- ELSE IF a static entry for the specific group address and the wildcard VID specifies Forward THEN Forward;
- ELSE IF a static entry for the specific group address and the wildcard VID specifies Filter THEN Filter;
- ELSE IF the Table 8-6 result for "All Group Addresses" is Registered THEN Forward;
- ELSE IF the Table 8-6 result for "All Unregistered Group Addresses" is Registered THEN Forward;
- ELSE IF a dynamic (Group Registration) entry for the specific group address and the frame's Base VID specifies Forward THEN Forward;
- ELSE Filter.

~~Where a given VID is allocated to the same FID as one or more other VIDs, it is an implementation option as to whether~~

~~a) The results shown in Table 8-7 directly determine the forwarding/filtering decision for a given VID and group MAC Address (i.e., the operation of the Bridge with respect to group MAC Addresses ignores the allocation of VIDs to FIDs); or~~
~~b) The results for a given MAC Address and VID are combined with the corresponding results for that MAC Address for each other VID that is allocated to the same FID, so that if the Table 8-7 result is Forward in any one VLAN that shares that FID, then frames for that group MAC Address will be forwarded for all VLANs that share that FID (i.e., the operation of the Bridge with respect to group MAC Addresses takes account of the allocation of VIDs to FIDs).~~

~~NOTE 2—In case d), the implementation effectively operates a single FDB per VLAN for group MAC Addresses. In case e), the implementation combines static and registered information for group MAC Addresses in accordance with the VID to FID allocations currently in force, in much the same manner as for individual MAC Addresses.~~

When forwarding or filtering a frame with a destination group MAC Address, a VLAN-aware Bridge may:

a) Ignore the allocation of VIDs to FID, and use Table 8-7 directly for the frame's Base VID; or
b) Take the same decision for all VIDs allocated to any given FID, forwarding if Table 8-7 specifies Forward for any VID allocated to the same FID as the frame's Base VID, and filtering otherwise.

Change Table 8-7 as follows:

Table 8-7—Forwarding or Filtering for specific group MAC Addresses

All Group Addresses control elements for this VID and Port specify (Table 8-6):	All Unregistered Group Addresses control elements for this VID and Port specify (Table 8-6):	Static Filtering Entry Control Element for this group MAC Address, VID and outbound Port specifies:				
		Registration Fixed (Forward)	Registration Forbidden (Filter)	Use Group Registration Information, or no Static Filtering Entry present. Group Registration Entry Control Element for this group MAC Address, VID and outbound Port specifies:		
				Registered (Forward)	Not Registered (Filter)	No Group Registration Entry present
Not Registered	Not Registered	Forward	Filter	Forward	Filter (Filter Unregistered Groups)	Filter (Filter Unregistered Groups)
	Registered	Forward	Filter	Forward	~~Filter~~ Forward (Forward Unregistered Groups)	Forward (Forward Unregistered Groups)
Registered		Forward	Filter	Forward (Forward All Groups)	Forward (Forward All Groups)	Forward (Forward All Groups)

<<The corrected entry was clearly a bug. A Group Registration Entry of Not Registered can be created due to a registration on another port, and the removal of a Group Registration Entry that has Not Registered on all ports is implementation dependent.>>