Annex B (informative)

<<The existing Annex B not only refers in detail to the (now) unnecessarily complex VLAN learning constraints that have been removed by this amendment changes to 8.8 (as required to make the operation of SPB clear), but is also incorrect in a number of other respects. It appears that it was not updated when MSTP was added to 802.1Q, so the following replacement is long overdue.>>

Replace the text of Annex B with the following:

Shared and Independent VLAN Learning

Bridges can learn MAC Addresses from each received frame's source address field, so that subsequently forwarded frames whose destination addresses have been learned can be filtered to restrict their transmission to the LANs necessary to reach their destination. Addresses learnt from frames with one VLAN Identifier (VID) may or may not be used to filter frames with another VID, depending on the capabilities of the bridge implementation and management controls. If learned information is shared for two or more given VIDs, those VIDs map to a single Filtering Identifier (FID) and the term Shared VLAN Learning is used to describe their relationship. If the information is not shared, the VIDs map to different FIDs and Independent VLAN Learning is being used.

This annex discusses when independent learning is required, when shared learning is required, and when individual bridges in the same bridged network can use either for given VIDs. The following general rules apply to the mapping of VIDs to FIDs:

a) When two distinct stations with the same MAC Address are attached to different LANs in the same network, any VID used to transmit or receive frames to or from the first has to map to a different FID from any VID used by the second. A single system with two LAN interfaces behaves as two stations.
b) A VID that is allocated to the CIST, or an MSTI, or an SPT Set, has to map to a different FID from any VID allocate to another of those active topologies.
c) If an intermediate system in the bridged network receives frames with one VID and transmits them with another, without changing the source MAC Address and using active topologies that allow the transmitted frame to traverse any of the same bridges and LANs as the received frame, then the two VIDs have to map to different FIDs.
d) Each ESP-VID maps to a FID that is not used by any other VID.
e) If a number of VIDs are used to support the same VLAN, with bi-directional conversations between pairs of end-stations using one VID in one direction and the other in another, then the VIDs have to map to the same FID. The Shortest Path VIDs (SPVIDs) used by SPB (3.14) to support an SPT Set (3.10) provide one example of such a configuration. Another example is provided by an 'asymmetric' VLAN configuration in which a server transmits frames that have one VID assigned while frames from the server's clients are assigned to another. Each Bridge Port's VLAN membership configuration is used to allow frames from the clients to reach the server and frames from the server to reach the clients, while preventing direct communication between the clients.

If rules a) through c) are ignored delivery of frames to a station can be unpredictable, as that station's location can be learned inappropriately and can appear to change frequently. If rule c) is ignored learning from frames can interfere with the configuration of ESPs, or one ESP's configuration can be unexpectedly affected by another's (depending on the bridge implementation, see 8.8.8). If rule e) is ignored learning will not be effective and frames will be flooded continuously, wasting bandwidth in the network.

When none of the above rules, a) through e), apply VIDs can be mapped to the same or different FIDs, though mapping to the same FID can speed learning after network reconfiguration.

VLAN-aware Bridges implement as many FIDs as VIDs, and in the absence of any other configuration each VID maps to the FID with the same value. The MST Configuration Digest (13.7) ensures that neighbouring

bridges are using each VID in conjunction with the same active topology thus ensuring consistency for learned information, or forces the use of a single CST.

The Filtering Database of a VLAN-unaware MAC Bridge behaves as if a single FID were implemented, with all VIDs mapping to that FID. Such a bridge is not capable of supporting configurations that would require the use of rules a) through c), or the unconstrained configuration of ESPs as required by PBB-TE that would require the use of rule c). A VLAN-unaware MAC Bridge is also unable to support SPB or SPBB since those technologies continue to support the use of the CIST as well as SPTs, to ensure that the mere existence of connectivity is minimally dependent on configuration, and therefore require at least one use of rule b).