

Using VLAN Tags to convey priority

Mick Seaman, P802.1Q-Rev Editor

This note is part of a marked up copy of P802.1Q-Rev/D1.03 (a work-in-progress draft created in the course of P802.1Q-Rev/D1.0 SA Ballot comment resolution). The point of the markup is to identify text relevant to ballot comment I-65. This note has been included in the marked up draft to provide clickable links to relevant text in the draft.

The markup identifies text relevant to the addition and removal of VLAN Tags as frames are forwarded by VLAN-aware Bridges and the use (or otherwise) of Priority-tagged frames, i.e. frames with a VLAN-Tag but with a value of 0 in the 12-bit VID field. It was produced by searching for and examining all instances of ‘VID’, ‘priority-tagged’ (with and without capitalization), ‘untagged’, ‘TPMR’, ‘trunk’, and ‘access port’ throughout the document, as well as reviewing clauses 6.8 through 6.20. I have highlighted in yellow (view in Acrobat Reader with ‘Comment’ selected to view the list of marked up locations). There would have been little point to simply marking all of the ~4,000 search hits, but I have attempted to select text that most clearly states, introduces, contradicts, or reinforces technical points. This note references the most cogent of the ~100 markups. I would be grateful to anyone who has the time to check the draft and identify any contradictory points I have missed.

Lists of relevant technical highlights¹ and recommendations follows (2 pages).

1. Technical highlights

1.1 Bridges do not transmit priority-tagged frames

This point is summarized by the NOTE on page 197:

“NOTE—As all incoming frames, including priority-tagged frames, are classified as belonging to a VLAN, the transmitting Port transmits VLAN-tagged frames or untagged frames. Hence, a station sending a priority-tagged frame via a Bridge will receive a response that is either VLAN-tagged or untagged, as described in 8.5.”

This is a “NOTE”, i.e. not normative text, so its function is to describe a consequence of the specification of bridge operation, not to add additional requirements (such as additional outbound frame filtering).² The consequence is a result of the fact that a C-VLAN Bridge associates a VID (vlan identifier) with each and every received frame: 7.6(a) page 193, 8.3 (a.3) page 199, and the VID value 0 is reserved and not a valid value for the PVID used as the VID for received frames that are not already VLAN-tagged or a potential translation for a received VID.³ On transmit (6.9.2, pg 162) a frame is either VLAN tagged with the non-zero VID value or untagged.

Other consequences of the normative text referenced is that a C-VLAN Bridge does not add an additional C-VLAN Tag to an already tagged frame (that is to say to a frame whose initial data octets compose a C-VLAN Tag). It can remove a C-VLAN Tag encoded in the initial octets of the data, but will not double tag. Equally a C-VLAN Bridge does not look beyond a C-VLAN Tag encoded in the initial data octets to search for other encoded protocol identifiers or tags.^{4,5} If a frame that is double tagged is forwarded by a network of C-VLAN Bridges, then it can

¹ Some of these, and additional points, were covered in <https://www.ieee802.org/1/files/public/docs2022/q-rev-seaman-vlan-tag-encoding-ppt-0122-v0.pdf> (presented in CRG discussion 2022-01-11). This markup provides detailed supporting references.

² There is a special case exception, detailed in 6.13 “Support of the ISS for attachment to a PBN” (page 175) where a Customer Bridge attached to a Provider Bridge transmits an S-tagged Priority-tagged frame to communicate a frame’s priority to a Provider Bridge (which doesn’t recognize C-tags). In this respect the relationship between the C-VLAN Bridge and the S-VLAN Provider Bridge is the same as that between an end-station and a C-VLAN Bridge. An EDE-M (see 802.1AE) attached to a PBN can also use 6.13. In other EDE types the corresponding interface is internal to the EDE, so an explicit tag is not required to communicate priority.

³ Table 9-2 page 274 states that the null VID value (0) “shall not be configured as a PVID or a member of a VID set, or configured in any FDB entry, or used in any Management operation.”. The very last of these prohibitions “used in any Management operation” is not true, as the 0 value is used as a special case, matching some specified selection of values or indicating ‘see elsewhere’ for a supporting operation.

⁴ Port-and-Protocol-based VLAN classification (6.12, pg 172) is only applied when a frame has not already been assigned to a VLAN. Stream identification as specified in 802.1CB (called out on pages 212 and pages 215 of P802.1Q-Rev/D1.03) can examine subsequent identifiers, but does not change the contents of the frame and the stream identification functions defined do not include matches for repeated tags (as far as I know).

⁵ Which allows layering of the protocols that use and control the network, with some protocols (e.g. LLDP) being capable of deployment at multiple sub-layers, each sub-layer offering the ISS to the layer above and the layered instances being scoped by addressing rather than by the assignment of different protocol identifiers.

eventually be transmitted double tagged, single tagged, or untagged.⁶ A frame that is, at any stage, single tagged or untagged will not be subsequently forwarded as a priority-tagged frame.

1.2 VLAN-unaware components do not remove VLAN tags

IEEE 802.1Q uses the layered protocol concepts introduced by the OSI Reference Model (6.1, page 149), i.e. layered communication occurs between peers. A station that adds a tag has one or more peers that remove that tag (or, equivalently, filter frames whose initial octets are those of the tag). The information in a tag can be observed, and used, by components that relay frames between the peers. A TPMR, for example, can use 6.20 “Support of the ISS with signaled priority”, page 186, for priority recognition but 6.20 does not remove the tag from the frame.

Tagging and detagging functions support the EISS (page 160).

1.3 VLAN-unaware end stations transmit and receive only untagged frames

As stated in 6.3 “Support of the MAC Service” (page 150) “When a VLAN-unaware end station is attached to an individual LAN, it will transmit and receive only untagged frames, and thus uses the VID(s) assigned by the Bridge Ports that classify those frames.”

A priority-tagged frame is not an “untagged frame” (3.286 on page 111).

Also in 6.4, on page 150, “NOTE 2-The LLC Entity in a VLAN-unaware end station can receive MAC Service indications that correspond to the receipt of VLAN-tagged frames but, being unable to recognize the tags EtherType value, will discard the frame and not deliver it to any LLC User.”

A problem?

Contrary to the statements immediately above, Clause 36.1.1 “Overview” in 36.1 “PFC operation” says “A VLAN unaware end station can use PFC by ending traffic as priority-tagged and by ignoring the VID in received frames” (pg 1567).

This text might be the (or one) cause of the issue presented in I-65 as it would be natural to configure a Bridge Port attached to a VLAN-unaware end station to transmit untagged frames. The Bridge Port’s detagging of a frame could then reveal an inner VLAN tag without a peer to peer relationship, with the possibility of the end station removing that inner tag to reveal a further protocol identifier that being recognized and acted upon.

A better end station model for PFC support

A more accurate approach would be to model the VLAN-unaware end station functionality within system attached to the Bridge Port as being supported by a VLAN-aware shim that provides the ISS. That shim would Priority-tag a frame on transmit to convey the ISS requested priority to the Bridge Port, and extract priority from and remove a received frames VLAN-tag (if present). If the Bridge Port is configured to (and can be trusted to) only transmit frames for a single VLAN to the attached end station’s system then there is no need for that system to be configure with the VLAN’s VID (or indeed to examine it on receipt).

If we standardized this behavior in the future we should do it properly and add a 6.n clause for the relevant behavior (a VLAN-aware shim extracting priority from a received VLAN-tagged frame and discarding the tag), simple though it might appear (this explicit approach of providing an explicit description is similar to the way we support TPMRs with 6.20). The PFC description is significantly looser than other parts of the standard in its adherence to the architecture and related issues have already be raised and are anticipated to be dealt with by the proposed PFC enhancement PAR.

A similar functionality is currently provided by the EISS Multiplex Entity (6.17, page 181) that allows VLAN-unaware clients of the ISS to use individual SAPs that are multiplexed over the EISS. However the 6.17 description is heavily oriented towards the support of CFM shims in interface stacks and does not included detail.

Caution

When mapping any provision of 802.1Q to vendor specific features or documentation, note that the term “trunk” is not used in any normative text to describe VLAN reception, transmission, or forwarding behavior. “trunk” is only used in Annex F (informative) F.1.3.2 “Rooted-Multipoint” (page 2066) to describe an application of asymmetric VLANs, and in Annex H (informative) “Interoperability considerations” [which discusses the need for consistent

⁶ Double tagging, though not in the standard, is used in some networks and regarded as a legitimate use of the control and data sub-layering of bridged networks (see [1.2](#)).

VLAN configuration and networks that comprise both VLAN-aware Bridges and VLAN-unaware MAC Bridges] (H.1.2, pages 2075, 2076, H3.2, page 2079, H.4.1, page 2080).

The term “access port” is used (capitalized to show that it is a ‘reserved term’ in a number of cases as in “Provider Access Port”) to identify ports on S-components in the PBN (Provider Bridged Network, Clauses 15 and 16) and EVB (Edge Virtual Bridge. The only other case in the normative text is on page 279 where it is used as part of a reference to 802.1X (Port-based Network Access Control). The term is used in Annex H (informative), H.2.1 “Consistency of static VLAN filtering”— “a common use will likely arose from managers who make use of an “access port construct”. An access port may be a port that is absent from the member set (8.8.10) for all VIDs but the untagged, default VID.”

Caution is therefore appropriate when configuring the network. Any informative description needs to be understood in terms of the provisions of the standard. An end station that needs to receive indications of frame priority in a VLAN Tag cannot be attached to an “access port” as informally described in H.2.1.

2. Recommendations

2.1 P802.1Q-Rev progression

P802.1Q-Rev/D1.0 has passed its SA ballot and we are therefore under an obligation to conclude the project and publish in a timely fashion. The document is arguably correct in the points discussed above, but potentially misleading in including “A VLAN unaware end station can use PFC by ending traffic as priority-tagged and by ignoring the VID in received frames” in 36.1. This statement should simply be removed. Supplementary material could be provided but the discussion of new material that is intended to clear up potential misinterpretation is usually a significant task, a subject of beauty contests and further claims of misinterpretation by casual readers. Exact replacement text is not clear at this time, which is a valid reason for not attempting its inclusion.

2.2 Future amendment P802.1Q-Rev amendment or revision

Clarification and correction of PFC material is clearly within the scope of the proposed PFC enhancement PAR (we are discussing existing functionality). This should take into account points raised above.

2.3 Network configuration

- 1) An 802.3 Bridge Port attached to end station that needs an explicit indication of the priority of received frames needs to transmit those frames VLAN-tagged even if those frames are for a single VLAN.
- 2) An 802.3 end station should be capable of receiving VLAN-tagged frames, even if its end station functionality is limited to use of a single VLAN.
- 3) An 802.3 Bridge Port attached to an end station should transmit all frames as VLAN-tagged.

Taken together (2) and (3) avoid the possibility of the Bridge removing a VLAN-Tag from a frame only for the end station to act upon a further VLAN Tag with a VID that should not (by network configuration) have been delivered to that particular end station. The alternative is for the administrator responsible for configuring the network to assume or trust that the end station is either incapable of recognizing a received VLAN Tag or has been configured to discard received VLAN Tagged frames. Since, for practical reasons, the administration of the end station or end station functionality is often the concern of a separate organization, this assumption or trust can be misplaced.