Ensuring that streams get transmitted tagged
Tony Jeffree

1. Summary

P802.1Qat/D5.0, which is the subject of the current Sponsor ballot, takes some liberties with the rules on where frames gets tagged, in order to enforce tagging of stream packets. Comment #26 on the Sponsor ballot (see summary at http://www.ieee802.org/1/files/private/at-drafts/d5/) from Steve Haddock offered three potential solutions, and various other solutions were discussed in the March 2010 Plenary and suvsequent AVB conference call.

The fundamental issue is that we want to be able to specify (1) SRP behaviour and (2) the defaults (in BA and otherwise) for Bridges, Talkers, and Listeners, such that we get "plug-and-play" behaviour of an AVB system; i.e., there should be no need for management intervention to get the plumbing to work. In order to get to that "plug and play" behaviour, one of the things that needs to happen is that whatever frames are transmitted as part of streams have to be transmitted Tagged, so that the priority value is carried end-to-end. Hence, whatever VID is being used for a given stream has to be configured along the path from the Talker to the Listener(s) so that the outbound Ports of Bridges along the path are (a) in the member set for that VID and (b) configured to transmit Tagged (i.e., there are no static configuration entries anywhere along the path that define that VID to be Untagged on those outbound Ports.)

It turns out that this is achievable, mostly within the existing conformance requirements for a VLAN Bridge, with only minor extensions to the specification of SRP and possibly P802.1BA.

2. What does a minimally conformant VLAN Bridge look like w.r.t. Tagging and VLAN support?

A good place to start is to look at the conformance requirements for a VLAN-aware Bridge component (5.4 in 802.1Q - http://www.ieee802.org/1/files/private/q-edition-drafts/2009%20edition/802.1Q-2009-edition-d0-1.pdf). The relevant requirements seem to be:

- Each port has to support at least one "acceptable frame types" value (receive Tagged only, Untagged only, or All frames). If all three are supported then the default has to be "All frames" [5.4 f) in Q]. So, for AVB purposes, we could (in the BA profiles) constrain this to be "All frames".
- Each Port's PVID defaults to 1 [5.4 g) and 6.9].

- Port & Protocol support is not required (5.4.1).
- Support configuration of at least one VLAN whose untagged set includes that Port [5.4 g)]
- There must be support for at least one VLAN [5.4 j)]
- Ingress filtering is disabled (8.6.2).
- Crucially, MVRP is supported [5.4 i), 5.4.2].

This means (I believe) that a VLAN Bridge that supports one untagged VLAN, using a PVID of 1 on all Ports, with a single static VLAN registration entry indicating Untagged for VID 1 on all Ports for VID 1 and all Ports in the member set for VID 1, and with capacity to support N more VLAN Ids (where N is determined by the marketers) through MVRP configuration, would be conformant to the spec.

(The existing spec allows many other flavours too – for example, a Bridge that accepts only Tagged frames, and has no VLANs at all

configured "out of the box" is I believe conformant, as is a Bridge that accepts only untagged frames and supports exactly one VLAN, using the default PVID which it transmits untagged on all Ports - so we could use P802.1BA to close down the options so that the version described earlier is the one required for AVB.)

3. How does this help?

A conformant VLAN Bridge has to support MVRP. Given that we already have to support MRP in order for SRP to work, this isn't much of a hardship – and if you're claiming to be a conformant VLAN Bridge you should be doing it anyway. What MVRP is designed to give you is one aspect of what we need – the ability to configure the Bridge Ports that point away from the Talker and towards the Listener to be in the member set for a given VID. So all that is required to make this happen is for the Listener to register for the VID concerned – and conveniently, the Talker Advertise in SRP tells the Listener which VID the stream will use.

So, the simplest version of the solution, using the Bridge config described in (2), goes as follows:

- The Talker advertises a stream, and in the appropriate field in the Talker Advertise, it communicates the VID that will be used to carry the stream.
- The Talker advertise is propagated in the usual way, and one or more Listeners decide to sign up for the stream. At the same time as sending the response, the Listener(s) also register for the VID associated with the stream, using MVRP[1]. This establishes the path from Talker to Listener for that VID.
- The Talker sends the stream, using the VID.
- The Talker **never** uses "0" as stream VID value; this would be regarded by SRP as a

misconfigured Talker Advertise and would be discarded.

Note that all VLANs configured in this way, using only GVRP, will be configured to be Tagged; it is only static VLAN registration entries that can create Untagged vlans, so if there is no other (than for VID=1) static VLAN registration entry, everything else is Tagged.

The beauty of this approach is that it requires no a-priori agreement among Talkers, Listeners, and Bridges as to what VID will be used for a given stream, it allows each Talker to use a different VID if it feels so moved, or even a different VID for each stream, and it uses stuff that we are (or should be) already using, and not only that, but we are using it for the purpose for which it was designed (configuring VLANs on demand).

The only downside seems to be that there could be a need to support a potentially large number of VIDs even for a consumer network if the Talkers all chose a VID per stream; this could be ameliorated for consumer devices by choosing a default stream VID value (any value other than VID 1) that would be used as the default stream VID by simple Talkers like the proverbial networked mic. This could be documented either in P802.1Qat or in P802.1BA.

4. What if someone does management?

If a manager chooses to add static VLAN configuration entries, then there is the potential that the VID used for a stream is statically configured somewhere on the path from T to L as Untagged. This is clearly broken as far as stream support is concerned, so SRP must detect that case as the Talker Advertise is propagated, and convert it to Talker Fail on an outbound Port if it detects that the VID associated with the stream is configured to be Untagged. There is no other implication, as there might be many more VIDs that are not mis-configured and that can carry streams successfully, so there is no impact on the domain boundary detection logic. So, as always with management, the manager has sufficient rope to hang himself

---

[1] To support this, the Listener needs only the "Applicant-only" state machine, and on a point-to-point link, only the point-to-point subset thereof.

with, and if he succeeds, that is his problem, but the system behaves correctly.

5. Alternatives to a fixed default VID?

There is the potential problem that whatever default VID we choose is already in use as the default for some other purpose, or that in a managed environment, we actually want to choose different stream VIDs for different Talker stations (for whatever reason), which militates against using a hard-wired value for the default stream VID. An alternative approach could work as follows:

- Each Bridge Port that supports SRP has a "default stream VID" parameter that is set to some chosen (and documented in Qat or BA) default value out of the box, but which is management read/write and so can be changed per-Port if need be.
- In the SRP domain boundary detection packet, we add a field that carries the default stream VID for the Port from which the packet was transmitted.
- The VID field in the SRP domain boundary detection packet is ignored by Bridge Ports and by Listeners, but is used by Talkers to determine the default VID value that they use as a stream VID if they are too dumb to do anything else.

This would allow the exact same behaviour as described in (3) if the defaults were unmodified, but would allow a network to be configured via management such that simple Talkers need not all use the same Stream VID if that was desirable.

6. Conclusions

It seems feasible to fix this problem with minimal changes to what we have currently specified in Qat, using mostly behaviour that is already specified for a conformant VLAN Bridge. If we can live with a fixed default stream VID, then the impact on Qat is very small – the only real change is changing Talker Advertise to Talker Fail if the VID is configured Untagged – and the only other change is the need for the Listener to register the VID for a stream using MVRP. If the

flexibility of configuring the stream VID for simple Talkers is desirable, then the additional VID field in the domain detection PDUs is needed, plus the ability to configure the default VID value per Bridge Port via management. Either way, the implications from a complexity standpoint don't look to be serious. The implications for BA, apart from specifying what we need from a minimal Q Bridge (which we had already agreed to do anyway), would be to describe somewhere in Clause 6 why this works.