

VLAN sub-interface YANG model

[draft-ietf-netmod-intf-ext-yang](#) &
[draft-wilton-netmod-intf-vlan-yang](#)

Rob Wilton (Cisco)

rwilton@cisco.com

IETF-IEEE coordination meeting, 6th June 2016

Proposed VLAN sub-interface YANG model

- Defines how to demux VLAN tagged traffic to independent IETF defined services (IPv4, IPv6, L3VPN, PWs, VPLS, EVPN)
- Many router vendors have proprietary configuration constructs similar to what is being proposed in these two drafts
- No standards exist for this technology in any standards body because historically the end user configuration has not been standardized
- However, there is now a strong market demand for automation via standard YANG models (c.f. OpenConfig)
- Without this draft (or equivalent), many IETF forwarding YANG models (as above) cannot interoperate with VLAN tagged traffic
- Members of NETMOD WG and IEEE 802.1 WG concerned with potential overlap with 802.1Q technology and associated YANG models

Comparison of IEEE and proposed YANG models

	IEEE 802.1Q bridge YANG model	Proposed Sub-interface VLAN YANG model
1. Standard modelled:	Model derived from IEEE 802.1Q abstract manageability model	YANG model is derived from core common subset of existing vendor models for VLAN classification and demux to sub-interfaces.
2. Forwarding paradigm:	Uses 802.1Q defined forwarding paradigm	Classification only, forwarding defined by separate IETF YANG models (IPv6, L3VPN, VPLS, etc)
3. Implemented on:	Implemented on 802.1Q bridge devices	Implemented on routers with support for L3VPNs, PWs, VPLS, etc.

802.1Q WG (TSN) – proposed next steps:

Concerns raised in 802.1Q WG presentation during 802.1 interim meeting - May 2016	Proposed resolution or next steps
1. IETF model must interoperate with IEEE 802.1Q	Update draft to add constraints on tag match and tag push ordering
2. IETF model may fundamentally violate IEEE 802.1Q architecture	Awaiting further IEEE 802.1 WG discussion
3. May overlap with 802.1Q host stack model	802.1Q doesn't define host stack mgmt interface. Awaiting IEEE 802.1 WG resolution
4. Flexible VLAN classification may put pressure on 802.1Q to implement the same features	Proposed IETF model has been used by many vendors for 10+ years, hence unlikely to be significant new pressures now
5. Only allow configuration that can be efficiently implemented	Proposed model is implemented in hardware by multiple device and ASIC vendors
6. May be better if draft is Informational/Experimental rather than Stds Track	Consult NETMOD WG chairs & AD

Backup slide -Full IEEE 802.1Q WG TSN meeting minutes

Hi All,

I've tried to write up my understanding of the comments and mood of the YANG presentation that I gave in TSN late Tuesday afternoon regarding my proposed ietf draft (<https://datatracker.ietf.org/doc/draft-wilton-netmod-intf-vlan-yang>). For reference the presented slides are at (<http://www.ieee802.org/1/files/public/docs2016/new-yang-wilton-vlan-classification-0516-v01.pdf>)

Sorry but I wasn't able to capture everything, so if I have misunderstood, misrepresented, or missed any of the key points then please correct or annotate any of my statements below. Comments and feedback is of course very welcome and desired to see if it possible to resolve the WG's concerns and to produce a standard interface/sub-interface VLAN YANG model that is acceptable to the 802.1 WG.

1) There is a very clear and strong feeling that any IETF VLAN YANG model must inter-operate cleanly with 802.1Q. I see that this is a valid concern and I think one that we should address:

1 a) One particularly strong concern is to ensure that tags are always matched/inserted/used in the right order. I.e. for a double tag match ensure that the other tag is an s-vlan tag, and the inner tag is a c-vlan tag. From several people I've spoken to, this seems to be their fundamental objection to the model and resolving this would potentially go a long way towards making this model acceptable. It should be straight forward to implement these constraints using YANG must statements, and I'll look at updating the model to reflect this.

2) There is a concern that the proposed model may potentially fundamentally violate the 802.1Q architecture:

2 a) In particular, it was pointed out at various times that a 802.1Q tag is not just a delimiter on traffic (as it is being used in some of the SP and subscriber related use cases) but represents protocol operations instead. I.e. there is a concern that VLAN tags are being misused. In the shorter term, I cannot see how this can really be addressed. One long term solution could be to define a separate "mux/demux tag type" using its own Ethertype. Its unclear to me whether there is any real desire to spend time defining this now, and the fact that it wouldn't inter-operate with 802.1Q bridges would seemingly make it less useful.

2 b) Mick Seaman doesn't seem to really agree with the sub-interface model and thinks that an alternative model would be better, although it is unclear exactly what that alternative model would look like. I'll try and discuss this further with Mick if he is available.

3) There didn't seem to be consensus as to whether 802.1Q formally defines a host stack for inter-operating with a 802.1Q bridge (specifically VLAN-aware end stations). There has been some opinion expressed that this may be missing from 802.1Q and that should possibly be added. If this was done, and included a manageability interface then this would potentially overlap with the L3 part of my proposed model.

4) A concern was expressed that allowing more flexible classification might put undesired pressure on 802.1Q to develop new features (e.g. as QinQ historically did). Given that the configuration supported by my proposed model is to support features and protocols that were standardized 10+ years ago I'm not sure that I see defining a common configuration model for those features now would add any significant additional pressure now.

5) A concern was expressed that the model should only allow for configurations that can be reasonably efficiently implemented in hardware. I believe that this is already the case - given that the model is already implemented in hardware by multiple device and silicon vendors.

6) There was a suggestion that the WG might be more comfortable if the draft was marked as experimental or informational. I personally don't think that marking it as informational would be appropriate, but possibly experimental could be used.

At the end of the meeting, I was thanked for attending and presenting at the IEEE interim and the 802.1 WG understood and appreciated the desire to produce a standard that is acceptable to the 802.1 WG. I would also like to thank the 802.1 WG for giving me the time and opportunity to present.

Thanks,
Rob