Email from Scott (27 July 2020)

From: stds-802-yang@LISTSERV.IEEE.ORG <stds-802-yang@LISTSERV.IEEE.ORG> **On Behalf Of** Scott Mansfield
Sent: Monday, July 27, 2020 9:29 AM
To: stds-802-yang@LISTSERV.IEEE.ORG
Cc: Rob Wilton (rwilton) <rwilton@cisco.com>
Subject: YANGsters call on 28 July

The call on 28 July is meant to discuss the IETF draft related to https://datatracker.ietf.org/doc/draft-ietf-i2rs-yang-l2-network-topology/

I haven't seen any conversation on this mailing list on the topic. In order to have a fruitful discussion, it is imperative that people take the time to review the draft.

We should (at least) point out the following…

The mac-address examples in the instance tree are not in IETF canonical form per https://github.com/YangModels/yang/blob/master/standard/ietf/RFC/ietf-yang-types%402013-07-15.yang

For completeness we should also point out the known issue that the mac-address definition in the ietf-yang-types module is different than the definition in the ieee802-types module.

The reference to 802.1ad should be replaced with 802.1Q-2018 and the Q-in-Q terminology is archaic and should be replaced with current terminology (provider bridging). Clause 16 of 802.1Q-2018 would be a better reference.

The reference to 802.1ah should be replace with 802.1Q-2018 Clause 26.

I think the IEEE references should be in the Normative section of the internet-draft, but that is a matter for debate.

The internet-draft is mostly a YANG module that has no augmentation of IEEE YANG. The only import is to import the typedef for VLAN. None of the bridge YANG or concepts is used. This might be ok, considering this is a network-level topology YANG module, but I'm interested in more IEEE experts weighing in on that topic.

Email from Don Fedyk (28 July 2020)

You asked for a review of draft-ietf-i2rs-yang-l2-network-topology-15

> Note the MAC-address definitions in the latest version have been changed to the IETF Model.

Here is my input:
I2RS L2 Topology:

From the draft-ietf-i2rs-yang-l2-network-topology:

> This document defines the YANG data model for Layer 2 (L2) network topologies by augmenting the generic network (Section 6.1 of [RFC8345]) and network topology (Section 6.2 of [RFC8345]) data models with L2-specific topology attributes. An example is provided in Appendix B.

Note this is a topology model - an abstraction of what you would find in a Router link state database

```
module: ietf-network-topology
   augment /nw:networks/nw:network:
     +--rw link* [link-id]
        +--rw link-id            link-id
        +--rw source
        |  +--rw source-node?   -> ../../../nw:node/node-id
        |  +--rw source-tp?     leafref
        +--rw destination
        |  +--rw dest-node?   -> ../../../nw:node/node-id
        |  +--rw dest-tp?     leafref
        +--rw supporting-link* [network-ref link-ref]
           +--rw network-ref
           |       -> ../../../nw:supporting-network/network-ref
           +--rw link-ref       leafref
   augment /nw:networks/nw:network/nw:node:
     +--rw termination-point* [tp-id]
        +--rw tp-id                              tp-id
        +--rw supporting-termination-point*
                [network-ref node-ref tp-ref]
           +--rw network-ref
           |       -> ../../../nw:supporting-node/network-ref
           +--rw node-ref
           |       -> ../../../nw:supporting-node/node-ref
           +--rw tp-ref         leafref
```

Also From RFC 8345

> This is a base model in YANG of links and nodes.
> Links and Nodes are a basis for describing routing systems.
> It is a layered system.
> L2 is an underlying network. (a VPLS/NVO3/VLAN/EVPN topology- LAN topology).

If we were to describe L2 Networks as a Nodes and Links Model - RSTP/MSTP/SPB would be the reference model.
We can describe as a links state database even thought the control plane may only see a distance vector summary.
The core is an active VLAN topology with nodes and links.
This not about what is on the wire IE MAC addresses an particular VIDs.

There is already a RFC 8346 for L3 Networks.

```
module: ietf-l3-unicast-topology
```

```
augment /nw:networks/nw:network/nw:network-types:
  +--rw l3-unicast-topology!
augment /nw:networks/nw:network:
  +--rw l3-topology-attributes
     +--rw name?   string
     +--rw flag*   l3-flag-type
augment /nw:networks/nw:network/nw:node:
  +--rw l3-node-attributes
     +--rw name?        inet:domain-name
     +--rw flag*        node-flag-type
     +--rw router-id*   rt-types:router-id
     +--rw prefix* [prefix]
        +--rw prefix    inet:ip-prefix
        +--rw metric?   uint32
        +--rw flag*     prefix-flag-type
augment /nw:networks/nw:network/nt:link:
  +--rw l3-link-attributes
     +--rw name?      string
     +--rw flag*      link-flag-type
     +--rw metric1?   uint64
     +--rw metric2?   uint64
augment /nw:networks/nw:network/nw:node/nt:termination-point:
  +--rw l3-termination-point-attributes
     +--rw (termination-point-type)?
        +--:(ip)
        |  +--rw ip-address*       inet:ip-address
        +--:(unnumbered)
        |  +--rw unnumbered-id?    uint32
        +--:(interface-name)
           +--rw interface-name?   string
```

We can see by this that it is Nodes , Links and termination points.
Nodes have a Name and prefix (destination) and metric (which is how far away from the location of the model
If the model was on a node).

Links have an identifier (numbered or unnumbered) and a link metric. (Single Hop metric)

Now let's look at what is modeled for I2RS L2 node:

```
augment /nw:networks/nw:network/nw:node:
  +--rw l2-node-attributes
     +--rw name?                  string
     +--rw description?           string
     +--rw management-address*    inet:ip-address
     +--rw sys-mac-address?       yang:mac-address
     +--rw management-vlan-id?    dot1q-types:vlanid {VLAN}?
     +--rw flags*                 node-flag-type
```

For a L2 Node We have a Name, an IP management address and a system MAC address.
They also have a VLAN ID

In IEEE the equivalent of a routerID is a Bridge ID or which is MAC address + a priority.
For SPB it is a System ID.
I'm OK with the mac address (which part of the bridge bridge ID or system ID is derived from) but I'm not 100% sure it has to be the same.
In IEEE there is a VLAN which is an active topology the ID it uses is not necessarily unique.

It would be better to have a BridgeID/SytemID and VLAN name.
A vid for the Management address supported on a VLAN is OK.
And there can be multiple nodes on a physical box. I'm not sure how L3 captures virtual topologies but IEEE supports virtual.

The simple case above is an interface with no L2 control plane component. That is possible but it is not an L2 node in my books.
Is it a termination point? I don't know.

Now let's look at I2RS L2 Links:

```
  augment /nw:networks/nw:network/nt:link:
    +--rw l2-link-attributes
       +--rw name?    string
       +--rw flags*   link-flag-type
       +--rw rate?    uint64
       +--rw delay?   uint32
```

In IEEE our links are really single or Link aggregation groups.
We also have metrics:

- Port path costs for Spanning tree
- Link metrics for SPB.

These are on a LAN adjacency basis.
Physical ports can share multiple LANs (using VLAN names) and each VLAN may support multiple VIDs.
All this is in IEEE802.1Q.
I think it should capture both links and LAG.
This assumes the possible multiple virtual networks that use the links all share the bandwidth resources there are attributes that do policing and shaping on various criteria.

Then there is the termination point. In the IP model it is the Link identifiers in IP format and interface name.
In I2RS L2 Proposal for Termination point:

```
augment /nw:networks/nw:network/nw:node/nt:termination-point:
    +--rw l2-termination-point-attributes
       +--rw description?                    string
       +--rw maximum-frame-size?             uint32
       +--rw (l2-termination-point-type)?
       |  +--:(ethernet)
       |  |  +--rw mac-address?              yang:mac-address
       |  |  +--rw eth-encapsulation?        identityref
       |  |  +--rw lag?                      boolean
       |  |  +--rw member-link-tp*           -> /nw:networks/network
                                             /node/nt:termination-
point/tp-id
       |  |  +--rw auto-negotiation?         boolean
       |  |  +--rw duplex?                   duplex-mode
       |  |  +--rw default-untagged-vlan?    dot1q-types:vlanid {VLAN}?
       |  |  +--rw vlans* [vlan-id] {VLAN}?
       |  |  |  +--rw vlan-id    dot1q-types:vlanid
       |  |  |  +--rw name?      string
       |  |  +--rw qinq* [svlan-id cvlan-id] {QinQ}?
       |  |  |  +--rw svlan-id    dot1q-types:vlanid
       |  |  |  +--rw cvlan-id    dot1q-types:vlanid
```

```
|   |  +--rw vxlan {VXLAN}?
|   |     +--rw vni-id?   vni
|   +--:(legacy)
|      +--rw layer-2-address?         yang:phys-address
|      +--rw encapsulation?           identityref
+--ro tp-state?                       identityref
```

I think termination points should be links with names and an identifier. The same as L3.
The identifier is not a MAC address it is a port identifier which is like an unnumbered link.
Or an unnumbered link in SPB.
A VLAN is a topology. IEEE does not give VLAN names they are defined as sets of supported VLAN identifiers. A frame belongs to a VLAN identified by its VID or lack of one.
What the model should use is a name. Identifying it by the VID set per TP would not bring clarity.
SPBV uses a Base VID, SPBM uses a Backbone VLAN (Identified by one or more VIDs) .
If the termination points enumerate the complete set of VLAN tags supported under a TP that should be a range of valid C-VIDs or SVIDs (with the CVIDs as wildcard).
The TPID should be used to identify the VLAN type. (C-VID or S-VID)
The VLAN IDs supported per VLAN is a Dynamic set where a control plane is used.
I'm not sure the set is necessary because it is merely informational IMHO.
The actual case of a S-VID with a particular C-VID is a special case where instead of the bridge components a bridge network is being used to provide a very specific VLAN termination point.
IEEE components look a each VID separately.
Auto negotiation and duplex are pure physical link properties. Do they belong here or at the link/member link level?

I cannot talk to the other L2 technologies.

These suggestions would be a much more apples to apples model with L3 network. That may not be what I2RS is looking at.
The other option I can see is a much simpler static termination point that ignores all the topology and control plane at layer 2.
However a true network topology model captures what L2 uses for network topology.

Cheers
Don

Email from Rob Wilton (29 July 2020)

Hi Scott, Don, all,

Don, first a big thank you for doing a detailed review and providing comments on this document.
In terms of the collective feedback that Yangsters provides the IETF with, if possible, I think that it would be useful to please categorise the feedback into two levels of importance:

(1) Comments which the Yangsters feel must be resolved (with discussion if necessary) before IETF publishes the document.

(2) Comments/recommendations that the Yangsters feel would improve the quality of document but are okay to leave to the authors/ADs/IETF's discretion as to how they proceed with them.

Regards,
Rob

Email from Don Fedyk (29 July 2020)

Hi Rob

Thanks It is a good question to ask as we review this. I think the IEEE position is based on what level of L2 topology capture is intended.

What I struggled with in the document is there are two main ways you can view an L2 network.
1) As a bridged topology (links and nodes (which are support bridge components)) that use bridging rules and spanning tree or shortest path trees for connectivity. (This is in line with how the L3 topology is captured. Nodes - links and metrics.)
2) As an interface to L2 where you get VLAN connectivity to another set of nodes over an opaque L2 Network.
To me this is a termination point - beyond that - the topology details are hidden.

The current document is more towards 2).
My recollection of I2RS is the topology is captured to a degree where forwarding decisions of routing can be understood by viewing the model. 1) above would do that for bridging as well.

Cheers
Don

Note to I2RS List

IEEE 802.1Q-2018 Clause 14.2.5 says that a Bridge Identifier is a 64-bit unsigned integer with the following parts:
Priority Component 4 bits
Locally assigned system ID extension 12 bits
Globally unique bridge address 48 bits