E-Tree Support in 802.1Q

Version 1

Stephen Haddock

January 21, 2010

802.1 Interim, Austin

Terminology

- Metro Ethernet Forum (MEF) defines three types of services, each based on an underlying type of Ethernet Virtual Connection (EVC):
  1. **E-Line:** based on a **Point-to-Point EVC**
  2. **E-LAN:** based on a **Multipoint-to-Multipoint EVC**
  3. **E-Tree:** based on a **Rooted-Multipoint EVC**
- Rooted-Multipoint (RMP):
  - Each end-point designated as either a Root or Leaf.
    - Roots receive frames from all other Roots and all Leaves.
    - Leaves receive frames from all Roots, but not from any other Leaves.
  - Similar to a what is commonly known as Point-to-Multipoint but allows multiple Roots.
  - Example application is to connect Subscribers to some kind of Content Provider or Service Provider.

Implementing E-Tree/RMP with VLANs

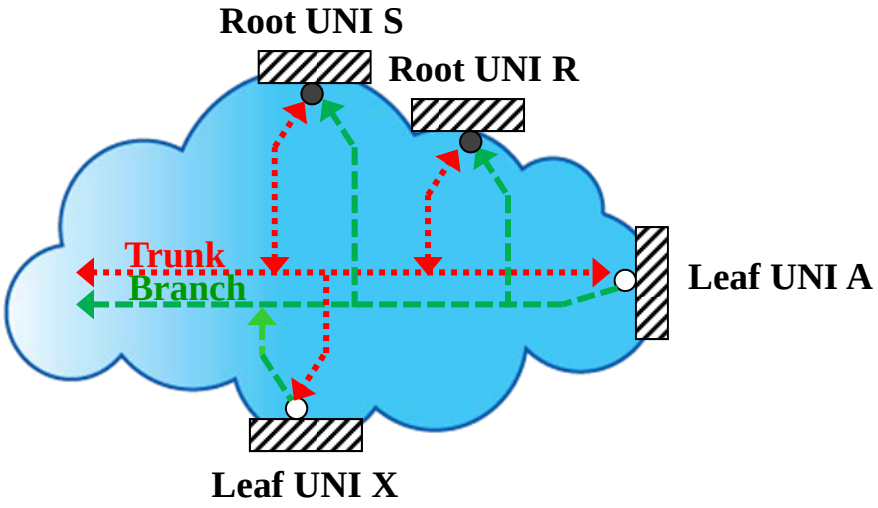

The RMP EVC is implemented with a pair of VLANs designated as a Trunk (red) and a Branch (green).

Based on the concept of Asymmetric VLANs (802.1Q annex B.1.3).

The Trunk and Branch VLANs are just ordinary VLANs except:

- Use Shared VLAN Learning on all bridges.
- Have specific filtering configuration at Root and Leaf ports:

At a Leaf:

- All ingress frames are assigned to the Branch VLAN.
- Egress frames on the Trunk VLAN are transmitted, egress frames on the Branch VLAN are filtered (not transmitted).

At a Root:

- All ingress frames are assigned to the Trunk VLAN.
- Egress frames on both the Trunk and Branch VLANs are transmitted.

Multiple network example

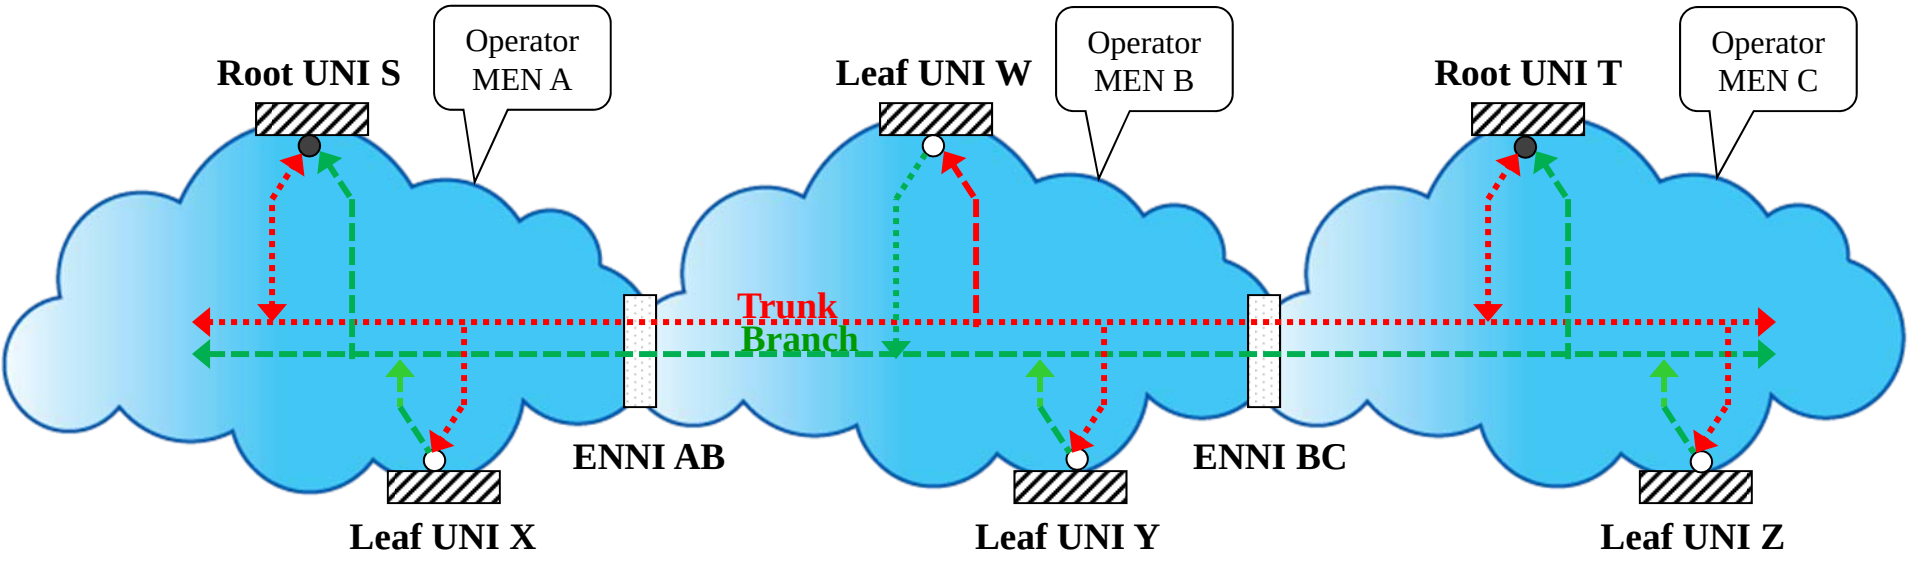

OVC End Point Subtypes: ● **Root** ○ **Leaf** ● **Trunk** ● **Branch**

The RMP EVC is implemented with a pair of VLANs designated as a Trunk (red) and a Branch (green). Both Trunk and Branch VLANs are carried across each ENNI.

Configuring RMP in 802.1Q

Leaf port configuration

- The Acceptable Frame Types parameter is set to the value *Admit Only Untagged and Priority-tagged frames*.
- The PVID is set to the value used for the Branch VLAN.
- The port is included in the Member Set and Untagged Set of the Trunk VLAN, but is not included in the Member Set of the Branch VLAN. This allows frames on the Trunk VLAN, but not the Branch VLAN, to be transmitted out the port.
- The Enable Ingress Filtering parameter is reset (i.e. disabling ingress filtering). This prevents frames received at the port from being discarded because the port is not in the Member Set of the Branch VLAN.

Root port configuration

- The Acceptable Frame Types parameter is set to the value *Admit Only Untagged and Priority-tagged frames*.
- The PVID is set to the value used for the Trunk VLAN.
- The port is included in the Member Set and Untagged Set of both the Trunk VLAN and the Branch VLAN.

See: http://www.ieee802.org/1/files/public/docs2005/qrev-haddock-p2mp%20vlan%20topologies-0105.pdf

E-Tree on Provider Networks

- All necessary capabilities and management controls to create E-Tree/Rooted-Multipoint services are supported at Port-based service interfaces to a 802.1ad Provider Bridged Network.
- Not fully supported at C-tagged service interfaces:
  - At a Root: A Provider Edge Bridge does not have the management controls necessary to map two S-VIDs to a C-VID (i.e. include the S-component CNP in the Member Set of two S-VLANs).
  - At a Leaf: A Provider Edge Bridge does not have the management controls necessary to have the S-component CNP in the Member Set of an S-VLAN with a different S-VID than the PVID value.
  - Stay tuned: suggested remedy in a few slides.

E-Tree on Provider Backbone Networks

- Can either multiplex both Trunk and Branch S-VLANs into a single backbone service instance (single I-SID), or map them two separate backbone service instances (two I-SIDs).
  - In either case, all RMP specific frame filtering is based on S-VIDs.
- Note that this approach works with any transport technology that can tunnel S-VLAN tags.

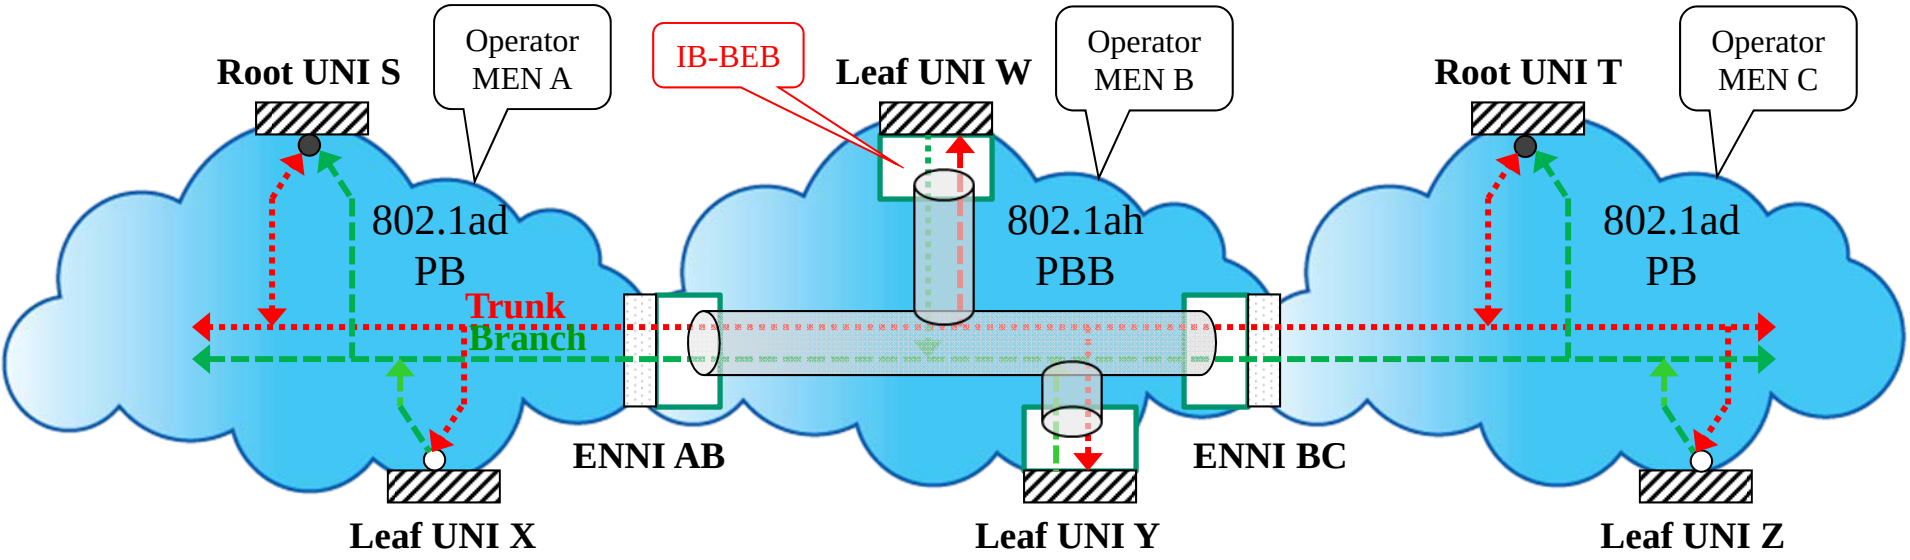

RMP VLANs and MVRP

- Since Trunk and Branch VLANs are just ordinary VLANs everywhere except at the Leaf and Root ports, MVRP can be used to propagate/prune Trunk and Branch VLANs through a network.
- Two caveats:
  - MVRP does not propagate Shared vs. Independent VLAN Learning information, so that needs to be configured in the bridges.
  - Since a Leaf Port is not in the Member Set of the Branch VLAN, the Branch VLAN must be declared on another port of the same bridge (e.g. a PNP on a Provider Bridge).

RMP VLANs and CFM

- CFM for Rooted Multipoint could monitor Branch and Tunnel VLANs individually, or map both VIDs to the same MEPs.
- Most likely these would be Up MEPs at Root ports and Leaf Ports.
- The relative position of Up MEPs and the “Port Filtering Entities” shown in Figure 22-4 means that each Leaf would see all other Leaves (as well as all Roots).
  - Not clear this is what is desired. May be preferred that a Root MEP sees all other Roots and Leaves, but a Leaf only sees Roots.
  - This would require making sure Egress Filtering (based on Member Set) would be enforced before CFM frames reached the MEP at the Leaf port.

Room for improvement

- The current Rooted Multipoint support in 802.1Q has two significant limitations:
  1. RMP frames must be untagged at Root and Leaf ports.
  2. The Provider Edge Bridge does not have the management objects to support RMP at C-tagged service interfaces.
- Both can be resolved by providing separate VID translation tables for ingress and egress.
  - This allows a 1:n VID mapping.
  - This is also required in .1aq Shortest Path Bridging (SPBV) to map between the Base-VID and SPVIDs

Separate Ingress/Egress VID translation

- Ingress and egress frames at Root and Leaf are tagged with a single RMP VLAN ID.
- In the network between Roots and Leaves, frames are tagged with Trunk or Branch VID (depending on whether frame entered network at Root or Leaf).
- In a Provider Edge Bridge, the RMP would be a C-VID while the Trunk and Branch would be S-VIDs

	Ingress Translation	Egress Translation
At Root	RMP → Trunk	Trunk → RMP Branch → RMP
At Leaf	RMP → Branch	Trunk → RMP (Branch discarded by Egress Filtering)

Possible Future Enhancements

- Restricted Flooding:
  - Currently frames sourced at a Root with unknown unicast destination addresses are flooded on the Trunk VLAN, potentially reaching all Leaves.
  - Could filter such frames at each Leaf (which assumes the address(es) of any device(s) attached at the Leaf port have been learned).
- Leaf Traffic Pruning:
  - Currently all Leaf-to-Leaf traffic is filtered at egress Leaf ports.
  - Is that a problem? Only frames that get that far are unknown floods and malicious attempts at Leaf-to-Leaf communication.
  - Once we have separate ingress/egress VID translation it is poosible to push this filtering deeper into the network by moving the RMP to Trunk/Branch translation point (to first bridge in from Leaf that would have Branch VLAN on more than two ports).
  - By management configuration or develop new protocol?

Thank You