



P802.1Qbh Draft 0.0 Introduction

Joe Pelissier

January, 2010

bh-pelissier-darft00intro-0110

Introduction

- P802.1Qbh specifies three major items:

- A Port Extender

- An M-Component which is used to make a Port Extender

- A “Extendable” VLAN Bridge (E-VLAN Bridge), a bridge that is capable of being extended using Port Extenders

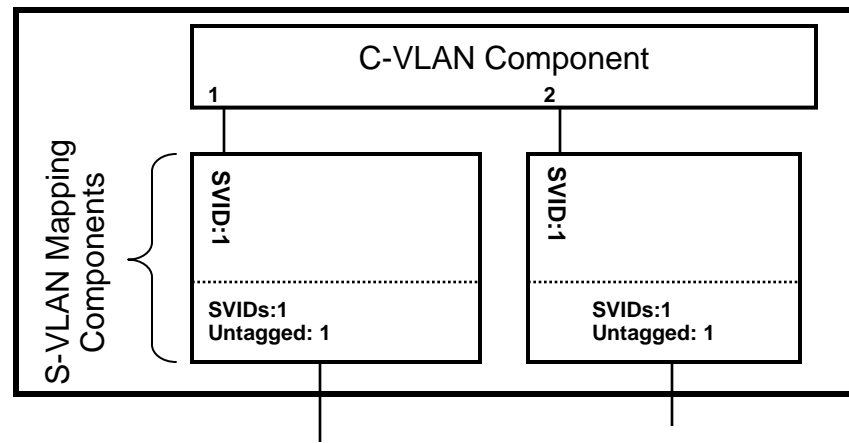
E-VLAN Bridge Model of Operation

- In the beginning...

For example, a two port bridge


SMC is transparent (mostly)

CVC ports 1 and 2 are not up since the SMC ports are not up



E-VLAN Bridge Model of Operation

- The Universe is created...
 - A Port Extender and an end device are attached
 - The SMC ports come up and thus the C-VLAN component ports 1 & 2 come up
 - LLDP starts executing on CVC ports 1&2
 - Port Extender discovered on port 1
 - Something else discovered on port 2
 - Port Extender Control Protocol starts executing on CVC Port 1
 - Configures member group and untagged member group on PE.



E-VLAN Bridge Model of Operation

- The Universe Expands...

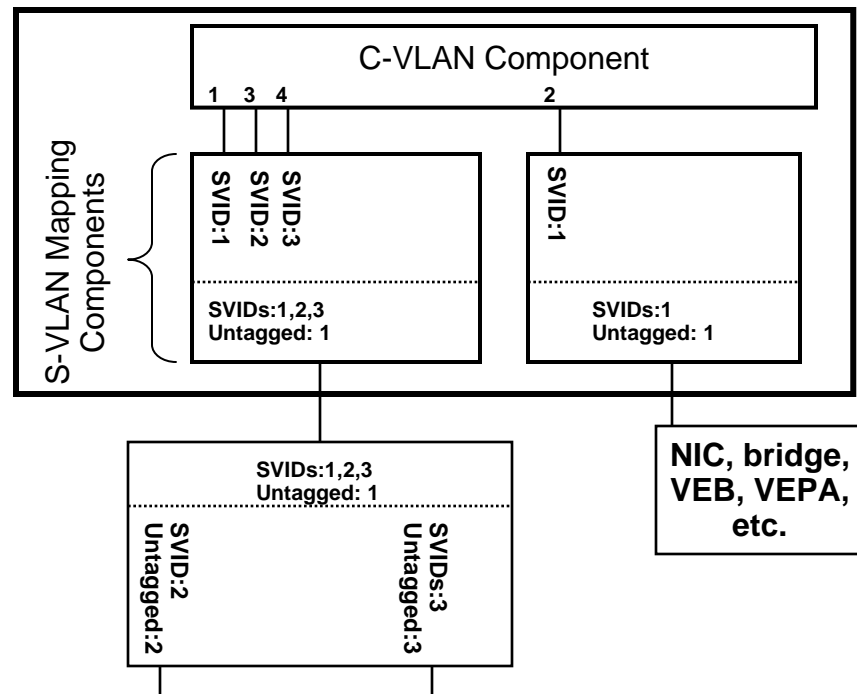
The E-VLAN Bridge discovers via the Port Extender Control Protocol that the PE has two ports active

CVC ports 3 & 4 instantiated

SMC SVID member groups and untagged member groups set-up

Using the Port Extender Protocol, the E-VLAN bridge programs the member groups and untagged sets in the PE.

The CVC starts executing LLDP on ports 3 & 4.



E-VLAN Bridge Model of Operation

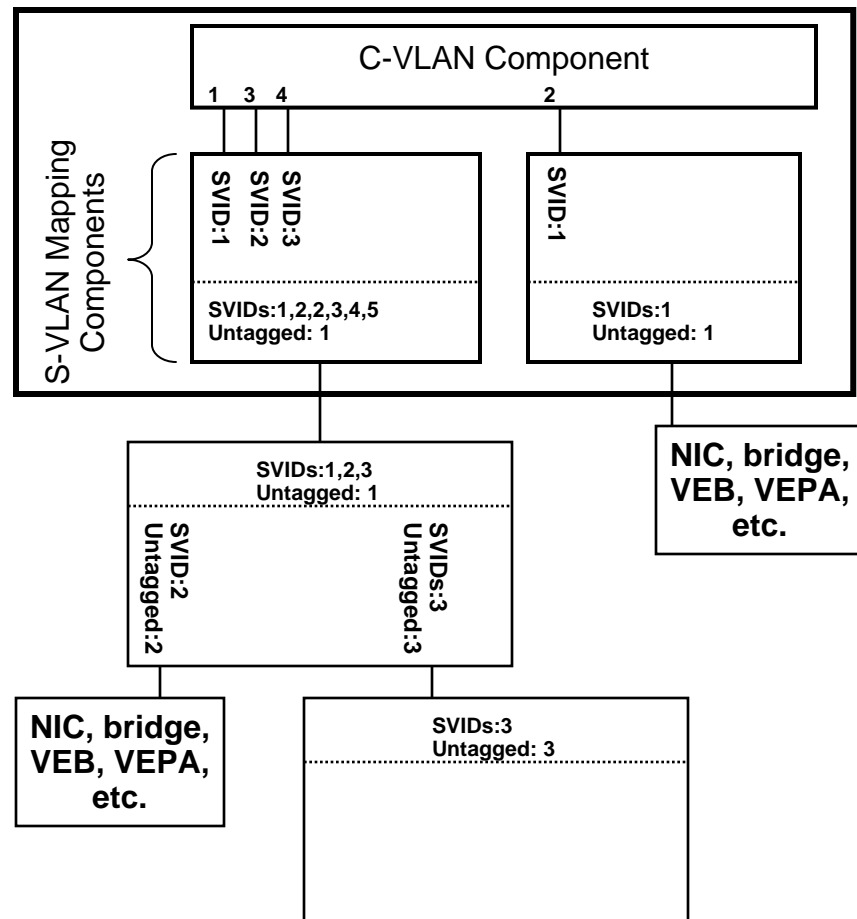
- The Universe Expands Some More...

LLDP on CVC port 3 discovers something other than a PE.

LLDP on CVC port 4 discovers the second PE.

Port Extender Control Protocol starts executing on CVC port 4

Configures second PE member groups and untagged member groups.



E-VLAN Bridge Model of Operation


- The Universe Expands Even More...

Port Extender Control Protocol on CVC port 3 discovers two end stations on the second PE.

CVC ports 5 & 6 instantiated

PE Control Protocol sets up member and untagged groups in 2nd PE

LLDP on CVC port 5 & 6 discovers something other than a PE.




E-VLAN Bridge Model of Operation

- The Universe Expands Just a Little Bit More...

(I need another PE to make my multicast examples interesting ;-)

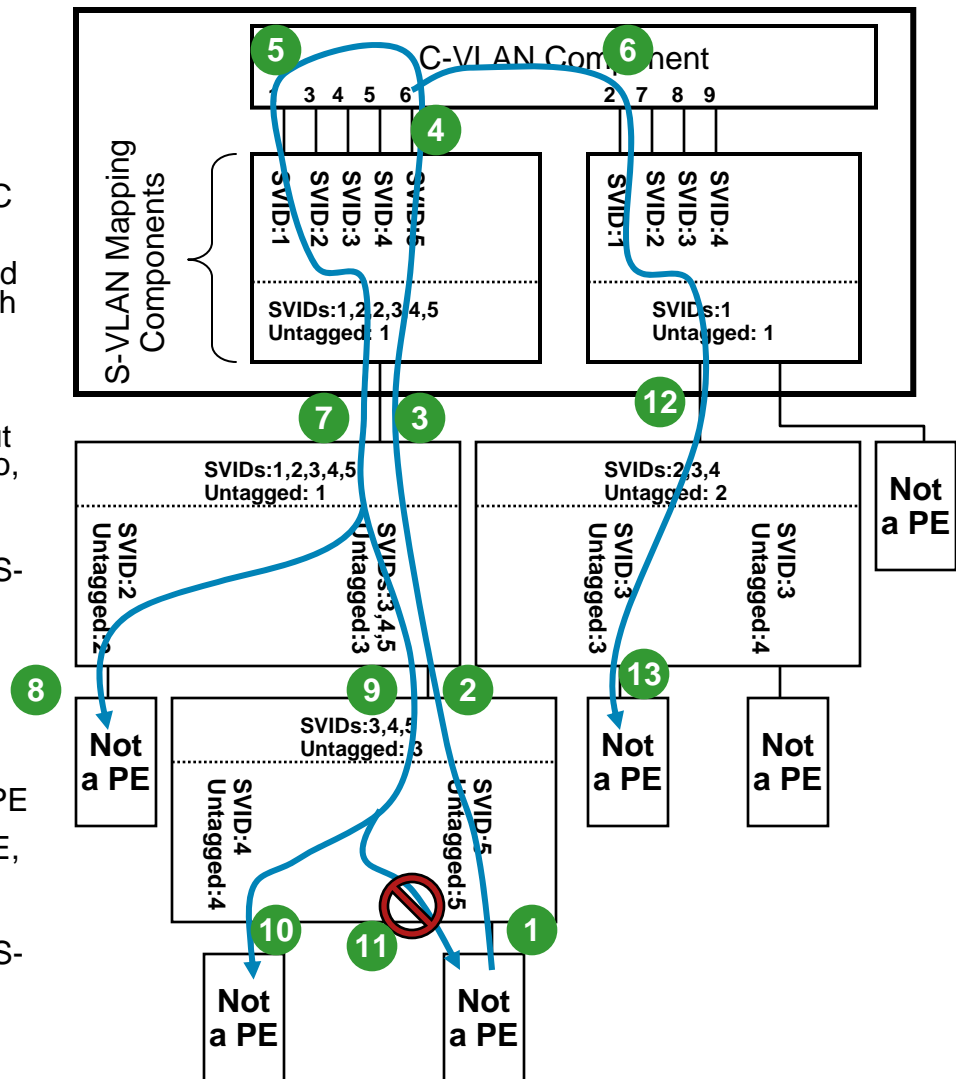
PE Control Protocol and LLDP operate as previously described

Details left as an exercise to the reader



E-VLAN Bridge Model of Operation (option 1)

- A Multicast Example...
 1. Multicast frame originated (no S-TAG)
 2. PE adds S-TAG with SVID 5
 3. SVID Still 5
 4. SVID removed, frame received on CVC port 6
 5. Without remote replication, frame would be forwarded on CVC ports 3 & 5. With remote replication, frame forward to Primary PE port with M-TAG; filter set true, source SVID set to 5.
 6. Frame could be sent to CVC port 8, but we'll keep this a multicast example. So, frame is M-Tagged with filter set false and SVID set to 0.
 7. Frame forwarded with M-TAG and no S-TAG
 8. PE does replication. M-TAG removed since this is the last PE.
 9. Frame has M-TAG and no S-TAG (implied SVID of 3)
 10. M-TAG removed since this is the last PE
 11. PE filters frame since this is the last PE, filter set TRUE, and SVID matches source SVID in M-TAG.
 12. Frame forwarded with M-TAG and no S-TAG
 13. M-TAG removed since last PE

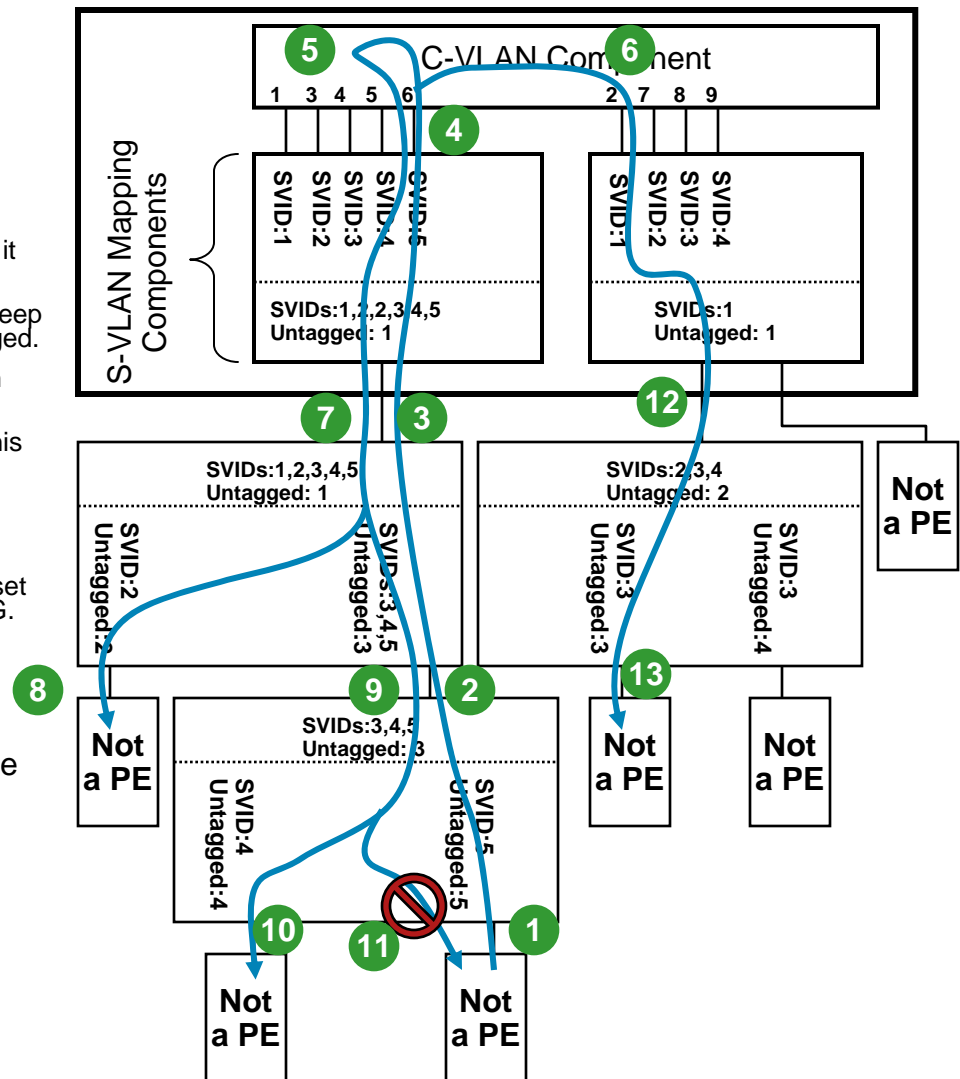


E-VLAN Bridge Model of Operation (option 2)

- A Multicast Example...

1. Multicast frame originated (no S-TAG)
2. PE adds S-TAG with SVID 5
3. SVID Still 5
4. SVID removed, frame received on CVC port 6
5. Without remote replication, frame would be forwarded on CVC ports 3 & 5. With remote replication, frame forward to the port on which it was received.
6. Frame could be sent to CVC port 8, but we'll keep this a multicast example. So, frame is M-Tagged.
7. Frame forwarded with M-TAG and S-TAG with SVID 5
8. PE does replication. M-TAG removed since this is the last PE.
9. Frame has M-TAG and S-TAG with SVID = 5
10. M-TAG removed since this is the last PE
11. PE filters frame since this is the last PE, filter set TRUE, and Port SVID matches SVID in S-TAG.
12. Frame forwarded with M-TAG and no S-TAG
13. M-TAG removed since last PE

Note: if a new PE was connected to the first PE, the first PE would remove the S-TAG and send the frame to the new PE.



Specification Issues

- Which base document to use?

 - PAR specifies 802.1Q-2005 (It was the only option)

 - 802.1Q-2009 will be current when this specification is done

 - Using the 2005 revision seems like a real pain for the next 802.1Q editor

 - 802.1Qbh is written to 802.1Q-2009 as the base document (maybe we need a PAR revision at some point?)


- E-VLAN Bridge model

 - Currently, the E-VLAN Bridge is modeled as a type of VLAN Bridge. This seems to the editor to be the most logical way to model it given that from a user's point of view its simply an enhancement. Alternatively, it could be modeled as a completely different type of bridge.


Specification Issues

- Port Extender / M-Component model

Currently, the M-component is modeled as a superset of the S-VLAN mapping component. Alternatively, the S-VLAN mapping functionality could be removed and a Port Extender could be modeled as a back-to-back combination of a S-VLAN mapping component and an M-Component. This makes the description quite awkward and is less similar to how these devices are actually likely to be built.



M-Component behaves as S-VLAN mapping component for non-M-Tagged frames. Performs replication for M-Tagged frames.



M-Component passes S-Tagged frames un-changed; makes multiple copies of M-Tagged frames with independent S-TAGs and SVIDS for each replication

Specification Issues

- Location of source SVID for multicast filtering

On multicast frames transmitted from the E-VLAN bridge, the SVID of the device that initiated the frame must be present to allow the final Port Extender egress port to filter the frame if it is the port that sourced the frame. The current text assumes this is part of the M-TAG. Alternatively, by clever selection of the C-VLAN component egress port, this may be indicated in a S-TAG prepended in front of the M-TAG. Placing this data in the M-TAG allows these frames to always be transmitted on the Replication Group Primary Port eliminating the need for the cleverness.

