# More Edge Bridge Spanning Tree

**Mick Seaman**

This note extends the description of VLAN-sensitive RSTP operation in ‘Provider Edge Bridge Spanning Tree’ to cover cases where a C-VLAN tag is allowed to select more then one S-VLAN. This can be used to provide E-LAN and E-TREE services over point-to-point S-VLANs.

## 1. Service instance selection by C-VLAN

The details of VLAN-sensitive RSTP operation in a Provider Edge Bridge†1 are very similar to the support of multiple trees by MSTP. The restriction that each C-VLAN only map to one S-VLAN reduces each of those trees to two ports, with the Customer Edge Port supporting multiple trees, and leads to the processing rules previously specified†2. This note considers how that restriction can be relaxed, and why.

Consider Figure 1. If C-VLANs from the customer’s LANs La, Lb, Lc are arbitrarily mapped to the three S-VLANs 1, 2, 3 within the provider’s network, then connectivity using one of them has to be blocked. This is easily avoided by requiring that each C-VLAN only mapping to one S-VLAN, as previously described.

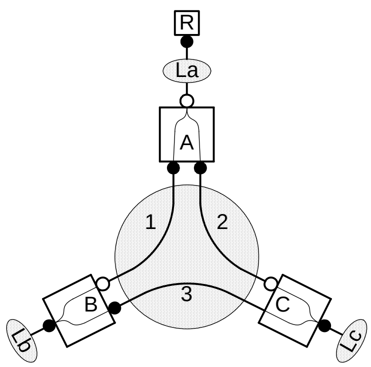


**Figure 1—Single service instance per C-VLAN**

An alternative is to require assignment of each C-VLANs to one of a number of disjoint sets, α, β, ... (say). Then the network might operate as shown in Figure 2.

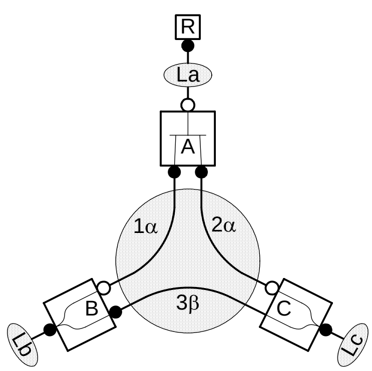


**Figure 2—Service instance flexibility**

C-VLANs α are being used by the customer equipment attached to La to communicate with both Lb and Lc. The spanning tree configuration remains the same as in the previous figure. However since both service instances 1 and 2 are carrying the same VLAN, the provider edge bridge C-VLAN component A now forwards between them. The potential loop is broken because instance 3 does not carry any of the C-VLANs α.

If there were any C-VLANs in common between 2 and 3, RSTP in bridge C would have to treat them as belonging to the same tree in the sense used by †1 to describe its relationship to MSTP. Service instances 2 and 3 would not then have been permitted to have distinct Rootward ports. The spanning tree topology shown in Figure 3 would have been the result, with 3 effectively unused.

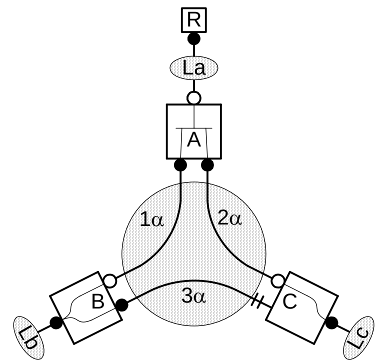


**Figure 3—Loop prevention**

Section 3 provides examples of networks that can be constructed, and the details of VLAN-sensitive RSTP operation are described in Section 4. However the first question to be answered is why bother?

## 2. Why

Clearly the network of Figure 2 requires the C-VLAN component ‘A’ to learn MAC Addresses. If the C-VLAN to S-VLAN mapping is under direct customer control, this may not be what the provider intended. It looks as if the customer is creating a multi-point service using the provider’s equipment after having arranged for point-to-point services. From this point of view the ‘correct answer’ to the problem apparently solved by Figure 2 is for the provider to provision a multi-point S-VLAN. Multicast and flooded transmissions through the provider’s network from A to both B and C will then be more efficient generally. A single frame need only be carried up to the point where the paths to B and C diverge.

However this plan assumes that the network provider actually has a plan to support multi-point S-VLANs, and has figured out to account for and explain the resources used so as to offer a viable service.There are doubts that this will be done in the near and medium-term. So Figure 2 actually looks quite an attractive way of starting to offer multipoint service. The real key is appreciating this is that a ‘service instance’ corresponds to an association between access points where a service is offered. A service instance doesn’t always have to be an S-VLAN, it may just happen

[†1] ./docs2005/ad-seaman-provider-edge-bridge-spanning-tree-0205-11.pdf
Thanks to the fact that the provider edge bridges are all neighbours, VSTP is just as functional (under reasonable constraints) as MSTP within an MSTP Region comprising just the Provider Edge Bridges.

[†2] In section 5 of ‘Provider Edge Bridge Spanning Tree’.

to make use of an S-VLAN. So configuration of the C-VLAN component ‘A’ may be entirely under the control of the provider, and S-VLANs 1 and 2 may not be exposed to the customer as being separate. While ‘A’ has to learn MAC Addresses there is now no requirement for an S-VLAN component to learn the same addresses, so no added burden has arisen.

Accounting for transmissions from La can be simply done in terms of A’s use of 1 and 2.

The configuration of such a service offering is not dissimilar to that of the popular frame relay offering where a head office serves a number of branches that mainly communicate with it rather than each other. It does not take advantage of the underlying multicast capabilities of the S-VLANs, but on the other hand reduces the requirement for intelligent assignment of the customer’s services to an S-VLAN with an appropriate topology.

Finally, the configuration of such a service does not preclude the simultaneous support of other connectivity, such as provided by S-VLAN 3 in Figure 2.

## 3. Examples

Figures 4 through 8 show how the connectivity of Figure 2 above should be affected by bridging paths external to the service provider and different locations for the customer’s Root Bridge.

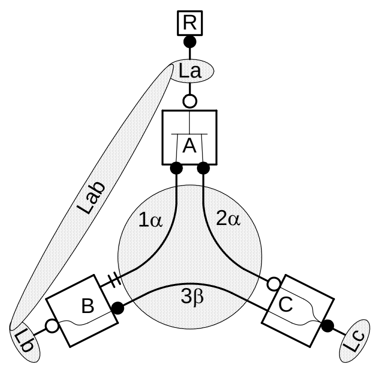


**Figure 4—An alternate path**

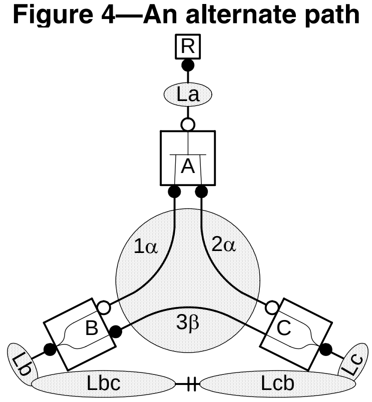


**Figure 5—An alternate path prevented**

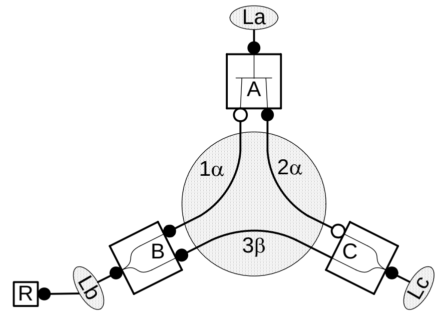


**Figure 6—Another Root Bridge location**

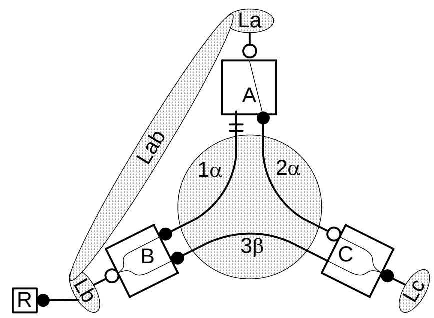


**Figure 7—With alternate path**

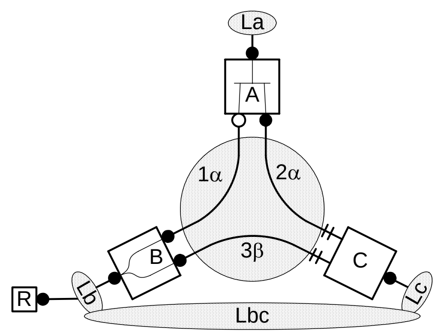


**Figure 8—With another alternate path**

## 4. Detailed specification

The following specification improves on that in †1. It works equally well for the more constrained case of assigning each C-VLAN to a single S-VLAN, and is suggested as the basis of P802.1ad specification for the Provider Edge Bridge C-VLAN component RSTP.

The specification uses the Provider Bridge Architecture summarized in Section A and is based on the RSTP state machines of 802.1D-2004 (as updated by the maintenance corrections that will appear in the P802.1Q-REV/D2.0).

The following procedure is added, and is executed whenever the assignment of C-VLANs to S-VLANs is changed:

1) updtTreePorts() assigns a tree number to each Provider Edge Port, each tree number representing a transitive closure of ports that have a C-VLAN in common, e.g. if port X has VLANs 1 and 2, port Y has VLANs 2 and 3, and port Z VLAN 3, then X, Y, and Z have the same assigned tree number†2.

The following changes are made to the existing procedures and state machine conditions:

2) The updtRolesTree() procedure (17.21.25) is modified to assign the Port Role of Rootward Port to the selectedRole variable of each of the Provider Edge Ports, fllowing the selection of the Root Port, iff:
   a) one of the Provider Edge Ports has been selected as the Root Port; and
   b) the Port would otherwise be assigned an Alternate Port Role; and
   c) no other Provider Edge Port with a better or equal spanning tree priority vector and the same tree number has already been assigned a Role of RootPort RootWard Port.

   If a Provider Edge Port with a worse priority vector has been assigned the RootwardPort Role, then tat port's role is changed to Alternate Port.
3) The global transition to the Port Role Transitions machine ROOT PORT state is extended to include a selectedRole of RootwardPort.
4) The procedures setReRootTree(), setSyncTree(), and setTcPropTree() each currently set a variable (reRoot, sync, and tcProp respectively) true for all (all other in the case of tcProp) ports of the Bridge. Each of these procedures is modified so that if the Root or Rootward Port in questions is:
   a) the Customer Edge Port, then the variable is set true for all (all other ports); but if it is
   b) a Provider Edge Port then the variable is set true for the Customer Edge Port and for all (all other) Provider Edge Ports with the same tree number.
5) The state machine condition allSynced currently requires synced to be true for all ports other than the Root Port. The definition of the condition is changed so that there is an independent value of allSynced for each Port of the Bridge (the condition is only used in the Port Role Transitions state machine (PRT) in the Root/Rootward and Alternate states), and if the Port is:
   a) the Customer Edge Port, then synced is true for all other ports;
   b) a Provider Edge Port, then synced is true for the Customer Edge Port and for all Provider Edge Ports with the same tree number.
6) Similarly, the definition of the state machine condition reRooted is changed so that there is an independent value for each Port of the Bridge and if the Port is:
   a) the Customer Edge Port, then rrWhile is zero for all other ports;
   b) a Provider Edge Port, then rrWhile is zero for the Customer Edge Port and for all Provider Edge Ports with the same tree number.

## A. Provider Edge Bridge Architecture

Customer equipment connected to a Provider Edge Bridge selects between S-VLANs by C-VLAN tagging transmitted frames. The operation of the Provider Edge Bridge is modeled as comprising two component bridges, as illustrated in Figure 9.

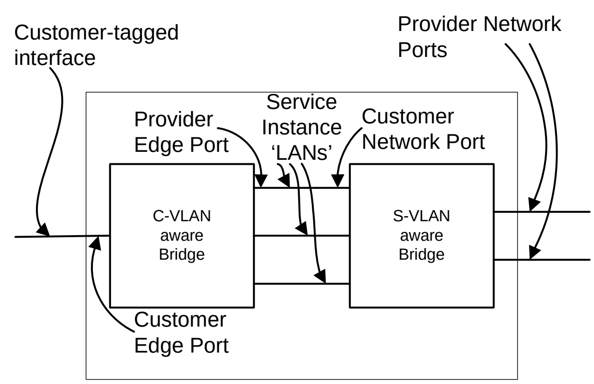


**Figure 9—Provider Edge Bridge**

The two bridges are connected by an number of internal LANs, one per S-VLAN. These are treated just like real LANs by each of the bridges, with a Bridge Port attaching to each.

---

†1 ./docs2005/ad-seaman-provider-edge-bridge-spanning-tree-0205-11.pdf.

†2 A simple algorithm, not necessarily the best, for doing this is to assign a tree number that is the same as the port number. Then each port is inspected, starting with the lowest port number, and the VLANs for all higher numbered ports with different tree numbers checked. If two ports with different tree numbers but which have a VLAN in common are found then all the ports with the higher tree number are changed to use the lower one.