Comment 5 Arjan de Heer

NAME: Arjan de Heer
COMMENT TYPE: T
CLAUSE: 5.9
PAGE: 13 LINE:
COMMENT START:
Each port of the Service-VLAN aware Bridge connects to either.. Connects to, or operates as?
Terminology is confusing: In the provider bridge there is a customer network port, but this does not necessarily connect to the customer, this is done by the provider edge port in some cases.
Names reflecting the function the port performs, these names are associated with the component:
* Selecting ports: S-VLAN aware ports that are connected to non S-VLAN aware components.
* Translating ports: S-VLAN aware ports that are connected to S-VLAN aware components that use a different VID assignment.
* 'Normal' ports: Otherwise
Names reflecting the place the port has in the network, these names are associated with the provider bridge (that may consist of more than one component):
* edge ports: Ports that are connected to VLAN aware components that have a different owner
* network ports: otherwise
These two names are orthogonal to each other. One is related to functions the other to the location in the network. There is no 1:1 mapping, it is not always the edge port that does the selection. This leaves a provider bridge with the following types of externally visible ports:
* network ports
* edge ports (in case of using a C-VLAN aware component port at the edge)
* selecting edge ports
* translating edge ports
Internally the provider bridge may have:
* selecting network port
* network port
The 4 external visible interface should are described in section 15.
* selecting edge port: port based
* edge port: customer tagged service interface
* translating edge port: provider tagged service interface
COMMENT END:
SUGGESTED CHANGES START:
Change naming as suggested
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 5

Discuss. It is still not clear to the editor what should be done with the proposed taxonomy, though the analysis is appreciated. In particular, labelling a port a "translating port" doesn't seem to have much meaning if it only translates one out of 4096 VLANs. I think the above taxonomy and the further complications that become apparent when additional connection scenarios are considered, point to the fact that no satisfactorily short port naming scheme will suffice to uniquely identify the ports function, role in the network or piece of equipment, and the ownership or control over the things that the port connects. It takes quite a bit of clause 15 and 16 to both identify all the interface properties and the service selection and segregation aspects. Attempts to build on Arjan's suggested scheme also run it to difficulties with the number of reserved words and words with loaded meanings. Using "service" and "interface" is fraught with difficulty. The terms "perimeter" and "boundary" might be acceptable, "border" is a little overloaded (BGP), as is "interior". The following diagram is intended to help the discussion by showing all the separate reference points that it might be useful to have names for (for simplicity provider-to-provider connections have been excluded), and discussing some of their properties. On the left of the diagram are items of equipment that are physically secured by the provider, a dotted line separates these from the items of equipment on the right which are vulnerable to customer or third party intrusion (intrusion into the equipment, not just into the connecting links). The diagram does not attempt to show who "owns" any item of equipment. S-VLAN aware components are labelled S, C-VLAN aware components are labelled C, and the label "cs" for "customer

system" is used to refer to any item of equipment solely under customer control that does not contain an S VLAN-aware component - i.e. it may be an end station, a VLAN-aware end station, a VLAN-aware or unaware router, a VLAN-aware (.1Q-2003) or VLAN-unaware Bridge. The round symbols denote LANs (possibly internal and microscopic).

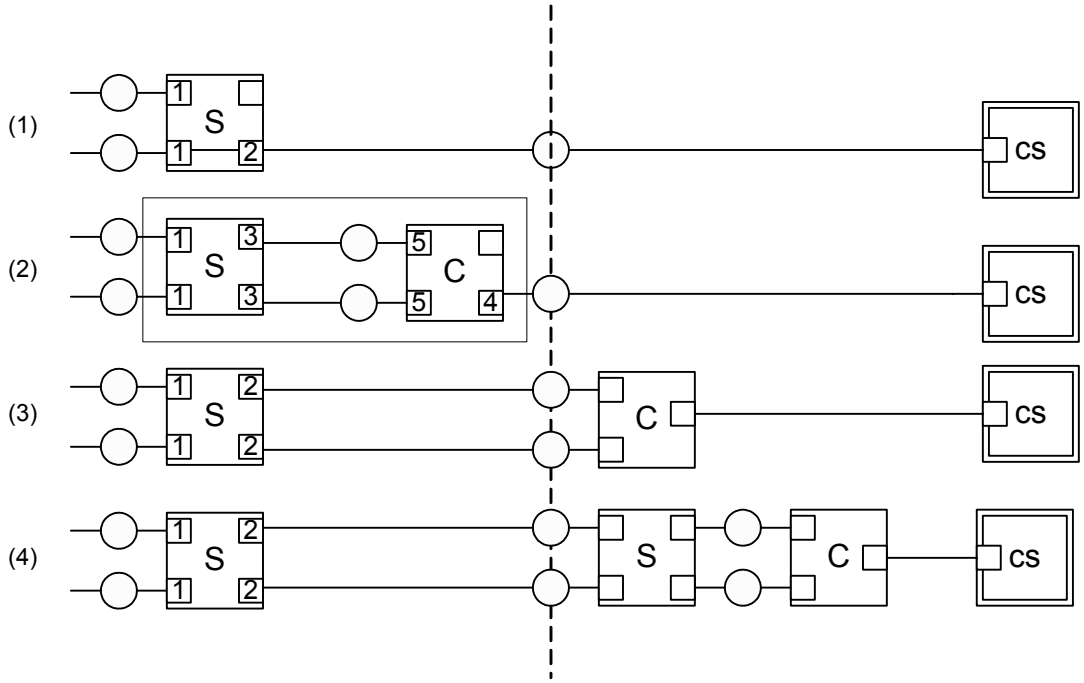

The ports labelled '1' all have the characteristic that they can carry traffic for any customer.

The ports labelled '2' all have the characteristic that they should only admit traffic to a service instance owned by the attached customer (the one who owns the 'cs' shown). They must not allow frames to ingress the network for any other service instance i.e. their P-VID must be set correctly and the combination of the S-VID member set, ingress rule checking, and translation table must not allow frames to ingress to an unauthorized service instance either.

The ports labelled '3' also have to perform this provider network ingress checking function, since the C-VLAN aware component in scenario (2) will not screen out S-tagged frames. It is a matter of supreme indifference to an attacker that he may be told not to send such frames. So the Ports labelled '3' are very much like those labelled '2' except that in any useful scenario they each receive frames, untagged, for only one service instance. And of course they are internal to a Provider Edge Bridge.

The port labelled '4' cannot only connect to a single customer, and the ports labelled '5', by virtue of being on the same C-VLAN aware component can only serve a single customer.

The above distinctions are quite independent as to whether a given provider decides that any of the ports labelled '2' will not provide service, because the customer should not S-tag frames. Whether the customer should or not, frames will have to be screened. So further configuration of the ports shown may perform further service instance selection, screening, and VID renumbering tricks, but as Arjan points out these functions are somewhat orthogonal. Further scenario (2) shows that screening (at ports '3') to ensure only the right traffic is admitted to a service instance, isn't necessarily a function of the first provider owned port (4) that customer traffic ingressing to the network encounters. So we could introduce more names to further reflect how ports are configured but I fear they would only result in a '2a', '2b' sort of distinction. The above labels being chosen to reflect the minimum necessary for correct service instance selection and segregation, not all things that might be desirable in a service offering.

Originally Ports ‘1’ and ‘2’ were called “Provider Network Port’ and ‘Customer Network Port’ which reflects my own experience in deploying scenarios (1) and (3). However it seems clear that a Port ‘4’ wants to be called a ‘Provider Edge Port’, at least for best fit with discussions in other forums. So we currently have:

Port 1 - Provider Network Port
Port 2 - Customer Network Port
Port 3 - ?
Port 4 - Provider Edge Port
Port 5 - ?

which looks decidedly odd when scenario (2) is labelled.

Avoiding the loaded terms “Service” and “Interface”, and trying not to overemphasize the provider’s point of view (the specification has to intelligible to customers as well) by avoiding the use of “Customer” to label a Port that is wholly within the provider network, the best suggestion I have so far is:

Port 1 - (Provider) Network Port
Port 2 - (Provider) Edge Port
Port 3 - (Provider) Edge Port (internal)
Port 4 - (Provider) Customer Edge Port
Port 5 - (Provider) Access Port

I’m not exactly excited about this labelling, but it seems better than “Reference Port 1”, “Reference Port 2”.

Discuss.