PROPOSED DISPOSITION OF BALLOT COMMENTS ON

IEEE P802.1ad/D2

Draft Standard for
Local and Metropolitan Area Networks—

Virtual Bridged Local Area Networks — Amendment 4: Provider Bridges

Sponsor
LAN/MAN Standards Committee
of the
IEEE Computer Society

Prepared by the Interworking Task Group of IEEE 802.1

Proposed Disposition of Ballot Comments prepared by:
Mick Seaman, Project Editor

This Proposed Disposition of Ballot Comments has been prepared to document the ballot comments received in the Working Group ballot on P802.1ad/D2.
The document contains:

1) A table of responses received.
2) A listing of comments received, each accompanied by a proposed disposition.

for discussion and resolution in the next 802.1 meeting.

Contents

1. Table of Responses

The following table indicates the status of each ballot response received in the P802.1ad/D2 ballot. Where comments have been received without an accompanying ballot, this is indicated in the ***Comments*** column.

As this was a Task Group ballot, as opposed to a full Working Group ballot, all participants in the activities of the 802.1 Interworking Task Group were entitled to submit a ballot, regardless of their voting status in 802.1. The voting status that is described in the following is therefore included only for information, and as an aid to planning future work. The table has been updated to reflect voting status as of 3rd February 2004.

The Status column indicates the voting status of the responder. ***V(oting)*** indicates 802.1 voting member at the start of the ballot period. **—** indicates a comment only response. ***L(iaison)*** indicates a voting liaison response. The ***Vote*** column indicates the vote cast; Y=Approve, N=Disapprove, T=Abstain due to lack of time, E=Abstain due to lack of expertise, O=Abstain for other reasons, C=Comments only.

Table 1-1—Table of responses

STATUS	VOTE	NAME	COMMENTS
V	T	Les Bell	—
V	N	Paul Bottorff	C
L (.3)	—	Richard Brand	—
V	—	Jim Burns	—
V	T	Marco Carugi	—
V	N	Dirceau Cavendish	C
—	C	Michael Chen	C
V	N	Paul Congdon	C
L (.11)	—	David Halasz	—
L (.11)	—	Russ Housley	—
V	Y	Arjan de Heer	C
V	Y	Anush Elangovan	—
V	—	Hesham Elbakoury	—
V	—	David Elie-Dit-Cosaque	—
V	Y	Norm W. Finn	—
V	Y	David Frattura	—
V	—	Gerard Goubert	—
V	Y	Steve Haddock	C
V	T	Atsushi Iwata	—
V	—	Neil Jarvis	—
V	—	Tony Jeffree	—
V	T	Manu Kaycee	—

Table 1-1—Table of responses

STATUS	VOTE	NAME	COMMENTS
V	—	Hal Keen	—
V	T	Bill Lane	—
V	T	Roger Lapuh	—
V	T	Loren Larsen	—
V	—	Yannick Le Goff	—
V	—	Marcus Leech	—
V	T	Mahalingam Mani	—
B	—	John Messenger	—
V	N	Dinesh Mohan	C
L	—	Tim Moore	—
V	—	Bob Moskowitz	—
V	—	Don O'Connor	—
V	—	Don Pannell	—
V, L(.3)	N	Glenn Parsons	C
V	—	Ken Patton	—
V	Y	Allyn Romanov	—
V	T	Dan Romascanu	—
V	Y	Jessy V Rouyer	—
V	Y	Ali Sajassi	C
V	—	Dolors Sala	—
V	Y	Mick Seaman	C
L (.3)	—	Matt Squire	—
L (.11)	—	Dorothy Stanley	—
V	N	Muneyoshi Suzuki	C
V	—	Jonathan Thatcher	—
V	—	Michel Thorsen	—
V	Y	Dennis Volpano	—
L (.11)	—	Jesse Walker	—
V	—	Karl Weber	—
V	T	Ludwig Winkel	—
N	—	Michael D. Wright	—

Table 1-2—Results (All Respondents)

CATEGORY	TOTAL	PERCENTAGE
Yes	10	62.5%
No	6	37.5%
Abs. Time	10	
Abs. Expertise	0	
Abs. Other	0	
No. of Voters	42	
Voters responding	25	60%

3. Ballot Comments

3.1 Scope

Comment 1 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: TR
CLAUSE: 1.1 PAGE: 3 LINE:
COMMENT START: It would seem appropriate to include a statement that this document also specifies a frame format for Service VLAN identifiers and priority.
COMMENT END:
SUGGESTED CHANGES START:
Add an item, perhaps around s) or t) that indicates a new VLAN frame format is defined for Service VLANs.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 1

Accept. Add an appropriate bullet to the list beginning with bullet (n). Ensure that this complements rather that contradicts existing bullet (f).

3.2 Terminology, Definitions, and Abbreviations

Comment 2 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: E
CLAUSE: 3
PAGE: 3-6 LINE:
COMMENT START:
clause 15 introduces a number of new terms that would benefit from having definitions in clause 3.
COMMENT END:
SUGGESTED CHANGES START:
Consider including the following; Customer Network Port, Provider Edge Port, Customer Tagged Interface.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 2

See disposition of comment 5. Editor to add place holder definitions for each of these, possibly renamed terms, to clause 3.

Comment 3 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: E
CLAUSE: 4
PAGE: 7 LINE:
COMMENT START:
Clause 16 introduces and uses a number of abbreviations that would benefit from having references in clause 4
COMMENT END:
SUGGESTED CHANGES START:
Consider including the following; CE, CPE, PPE, A-LAN, S-LAN, PIP, OAM&P
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 3

Accept.

Comment 4 **Paul Congdon**

NAME: Paul Congdon
COMMENT TYPE: ER
CLAUSE: 5.9
PAGE: 13 LINE:
COMMENT START:
I am confused by the statement that a customer network port can connect to customer system or a customer VLAN aware bridge. What is meant exactly by a customer system? It sounds like an end-station only, but I assume it could also be an old-fashion 802.1D bridge.
COMMENT END:
SUGGESTED CHANGES START:
Consider adding clarification of what a customer system is. Also, if the device is VLAN aware, does it mandate that it is customer VLAN aware (if there is any distinction between classic 802.1Q VLAN awareness and customer VLAN awareness).
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 4

Accept. The intent is that the customer system is customer controlled, and separate from the Provider Bridge. It could be an end station, bridge, or router, and may or may not be VLAN aware (of any type of VLAN). The distinction attempted here is between (a) a further component in the Provider Bridge which connects in turn to a customer system (b) connection directly to the customer system. The suggestion in the January meeting was that diagrams of the 'plan view' form used in Clause 15 be added here to explain the structure of the Provider Edge Bridge,. Hopefully that should help.

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).

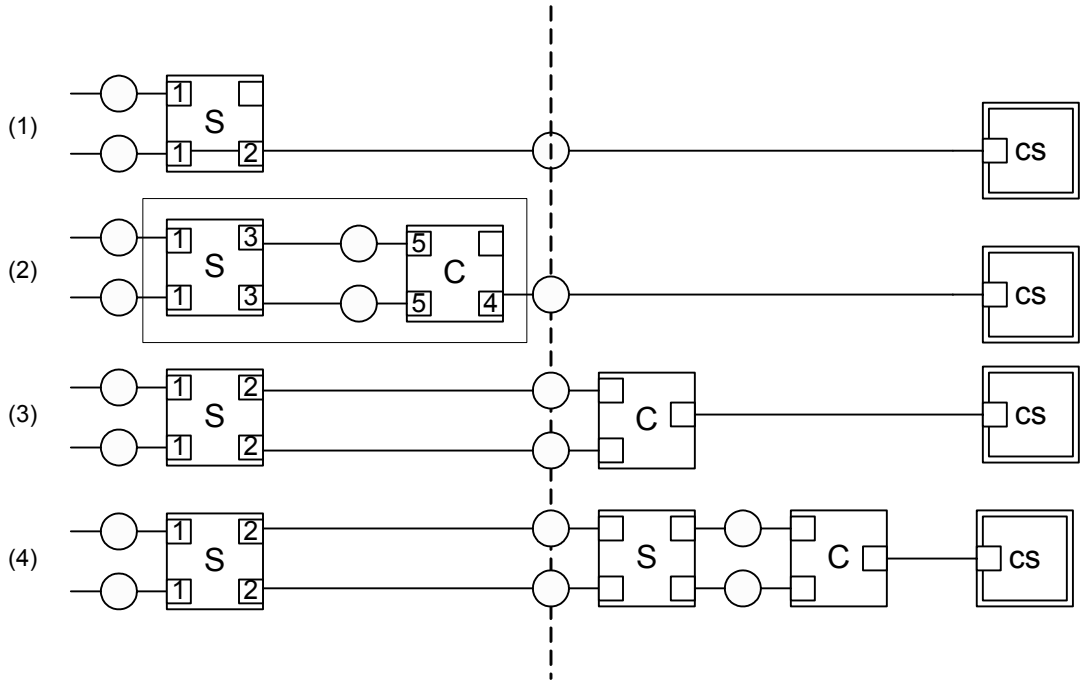

The ports labelled ‘1’ all have the characteristic that they can carry traffic for any customer.

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.

Comment 6 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: E
CLAUSE: 5.2
PAGE: 8 LINE:
COMMENT START:
The name "VLAN Bridge" is not much descriptive. It would be good to change it to "(Provider-Managed) Customer Bridge" through out the document since its primary purpose is to deal/ translate customer VLANs.
COMMENT END:
SUGGESTED CHANGES START:
d) Customer Bridge (5.8);
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 6

Reject. Though "VLAN Bridge" is not very descriptive it is used here deliberately so that there is continuity in terms from 802.1Q-2003. Remember the final output document is an amended .1Q, not just a Provider Bridges spec. There is always a risk that in specifying Provider Bridges we introduce unnecessary change for those who just want a Customer-VLAN aware Bridge, as 1Q-2003. Some change is clearly needed so as to maximize use of common ideas and components, rather than ending up with a subtly different VLAN capability for provider bridges, but this seems to be one point where terms can be carried forward. If this is not acceptable to the group, an alternative would be to define a new term as being a "VLAN Bridge as defined by 802.1Q-2003 and 802.1Q-1998", but the new term should not assume any interest on the part of the reader in Provider Bridges specifically.

Comment 7 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: E
CLAUSE: 5.8
PAGE: 13
LINE:
COMMENT START:
As mentioned previously, it would be good the change the title to " (Provider-Managed) Customer Bridge" to first be more descriptive and second to differentiate better between VLAN-aware Bridge Component and Customer Bridge.
COMMENT END:
SUGGESTED CHANGES START:
Customer Bridge Conformance
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 7

See resolution of comment 6.

3.3 Edge Ports and/or Edge Bridge

Comment 8 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: 5.2
PAGE: 8 LINE:
COMMENT START:
There is a special case of a Provider Bridge that includes a Service-VLAN aware Bridge component and one or more Customer-VLAN aware Bridge components. It is introduced in 5.9 as

a Provider Bridge that has Provider Edge Ports. I think the concept is significant enough that this type of system warrants a specific title, and creating such a title will help clarify the descriptions in section 5.9.
COMMENT END:
SUGGESTED CHANGES START:
Add bulleted item: "f) Provider Edge Bridge (5.9)."
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 8

As discussed at the January 2004 interim:

A "Provider Edge Bridge" should be a conformance class of equipment, as suggested here and in Steve's following comment.

Comment 9 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: 5.9
PAGE: 13 LINE:
COMMENT START:
Introduce the concept of a Provider Edge Bridge early
COMMENT END:
SUGGESTED CHANGES START:
Change the first paragraph of 5.9 to: "A Provider Bridge is a system that comprises a single Service-VLAN aware Bridge component. Each port of the Service-VLAN aware Bridge component operates as

a) A Provider Network Port; or
b) A Customer Network Port.

A Provider Edge Bridge is a special case of a Provider Bridge that further comprises one or more Customer-VLAN aware Bridge components. Each port of a Customer-VLAN aware Bridge component within a Provider Edge Bridge system operates as

c) A Provider Edge Port; or
d) An Internal Port.

Each Provider Network Port connects to a LAN in the provider network. Each Customer Network Port can connect either

1) Directly to a customer system; or to an
2) Internal Port of a Customer-VLAN aware Bridge component within a Provider Edge Bridge system.

Note xx: The connection between an Internal Port of a Customer-VLAN aware Bridge component and the Customer Network Port of the Service-VLAN aware Bridge component is made at the ISS interface of the respective ports. There is no need to have a MAC specific layers at this connection. Each Provider Edge Port connects directly to a customer system."
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 9

See resolution of comment 8.

Comment 10 Steve Haddock

NAME: Stephen Haddock

COMMENT TYPE: T
CLAUSE: 15.4
PAGE: 86 LINE:
COMMENT START:
<Details of comment>
COMMENT END:
SUGGESTED CHANGES START:
In Figure 15-4 change the "Provider Bridge" label of the box containing the S-VLAN aware and C-VLAN aware components to "Provider Edge Bridge". Change Figure 15-4 to show the connection between the S-VLAN aware and C-VLAN aware components to be at the ISS level of the Customer Network Port and the Internal Port.
Change the second paragraph of 15.4 to "The Customer-VLAN aware Bridge component assigns each frame received from the attached customer system to a Customer VLAN based on the C-VID or default PVID, and presents it to each Customer Network Port of the Service-VLAN aware Bridge component that supports a service instance that the customer desires to carry that VLAN. Similarly frames received by the Customer-VLAN aware Bridge component from the Service-VLAN aware Bridge component are assigned to a Customer VLAN and forwarded to one or more customer systems."
Change the third paragraph of 15.4 to "The Customer-VLAN aware Bridge component provides service instances for a single customer. The Service-VLAN aware Bridge component can provide port-based service instances to a number of Customer-VLAN aware Bridge components each serving different customers."
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 10

Relabel Provider Bridges as Provider Edge Bridges consistent with the proposed resolution of comment 8.

Reject the suggestion that there is such a thing as "connection at the level of the ISS". The ISS is a user/provider layer interface, not a peer to peer connection. Additionally such a connection would create a port to port connection without a LAN. Since the connectivity provided by spanning tree (and other protocols) is in terms of a bi-partite graph of bridge and LAN nodes with ports as the edges of the graph, such a port to port connection would prove an exception in every protocol description and proof - I'm not sure how two edges of a graph can connect directly together. What is needed here is a media independent 802.1 MAC/LAN/service (whatever name) for wiring up internal components in a system. This should be defined in a subclause of clause 6.5.

Accept the proposed text for the second and third paragraphs of 15.4. The fact that it specifies the result of the process, rather than guiding the reader through the detail of port to port transfer is fine. It should be understood from the figures that the C-VLAN aware component has ports internal to the Provider Bridge.

Comment 11 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 5.2
PAGE: 8
LINE:
COMMENT START:
The last sentence is not quite accurate since a Provider Bridge can NOT only include Customer-VLAN aware bridge component. Also the VLAN bridge (Customer bridge) can NOT just include Service-VLAN aware bridge component.
COMMENT END:
SUGGESTED CHANGES START:
and for the following systems
d) Customer Bridge (5.8) that includes instance of a Customer-VLAN aware Bridge component;

e) Provider Bridge (5.9) that includes instance of a Service-VLAN aware Bridge component and can include one or more instances of Customer-VLAN aware Bridge component.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 11

Reject. The proposed text just gets deeper into the problem of defining the systems, which will only get worse when the Provider Edge Bridge is added. The text at this point is simply trying to introduce the reader to the structure of the clause - core component, components based by picking from/adding to the core component functionality, systems comprised of these components. Suggest that replacing "than include instances of" with "that can include instances of" is sufficient wordsmithing to avoid the appearance of defining the systems at this point rather than in 5.8 and 5.9.

3.4 Priority signaling from Customer to Provider

Comment 12 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: 5.6.3 and 6.8 and 15.1
PAGE: 12 and 27 and 85 LINE:
COMMENT START:
Permitting customer bridges to use S-TAGs with null S-VID to signal a desired priority to a provider bridge is unnecessary and undesirable. It is unnecessary because the functionality can be achieved in a Provider Edge Bridge without having the customer-VLAN aware Bridge component create an S-TAG. It is undesirable because:

1) It creates a way for a customer to signal a desired priority to a provider that the provider must honor without the opportunity for the provider to re-map this to a different priority. (This is the case where the customer-VLAN aware Bridge component in a Customer Bridge system uses an S-TAG to signal a desired priority. A provider bridge then has no opportunity to re-map this priority since the User Priority Regeneration Table of 6.7.3 specifically excludes regenerating priority derived from tagged frames.) In practice this means that this portion of the Customer Bridge system would need to be manageable only by the provider, and not by the customer.

2) It violates a clean architectural concept that customer-VLAN aware Bridge components are only aware of Customer VLAN tags, and service-VLAN aware Bridge components are only aware of Service VLAN tags.

3) It violates the long-standing requirement that VLAN aware bridges do not transmit priority tagged frames. A better way to accomplish this functionality is to require that just as a C-Tag VID may be used as a service selector by connecting to a Provider Edge Bridge (containing one or more Customer-VLAN aware Bridge components), so a C-Tag priority may be used as a priority request by connecting to a Provider Edge Bridge. The priority at the ISS of the Internal Port will be taken from the C-Tag. The Customer Network Port connected at the ISS to the Internal Port thus receives a priority parameter on an untagged (no S-Tag) frame and may map that to a priority appropriate to the provider network using the User Priority Regeneration Table of 6.7.3. This has the further advantage that management control of the priority used on the provider network is now in the Service-VLAN aware Bridge component.
COMMENT END:

SUGGESTED CHANGES START:
Remove section 6.8. Remove sections 5.6.2 and 5.6.3. Remove bullet "f" of 5.6.1 and 5.7.1. Change the last sentence of 15.3 to "Frames transmitted through a Customer Network Port by a Customer-VLAN aware customer system do not include an S-TAG." In the first sentence of 15.4 change "allow service instance selection and identification by Customer VID (3.10)" to "allow priority selection and service instance selection and identification by Customer VLAN tag"
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 12

Discuss. This proposed disposition attempts to capture the discussion at the January meeting.

The comment is a complex one, a summary of the resolution so far is that the idea of recovering the priority from the C-TAG, as proposed in the comment, will be adopted as the most pragmatic way for a service provider to interface to existing customer equipment. How it is proposed that this be done is detailed below. However a number of the arguments made in the comment advocate the approach in the negative - i.e. by attacking the alternative. Some of these arguments don't hold up well, particularly as what the customer actually puts on the frames has to be viewed as out of the control of the service provider and must be policed by the provider in all cases. The interesting questions then become what needs to be done by the provider to check customer frames, with an emphasis on not introducing egregious extra mechanisms, and finding potential value in existing mechanisms.

Clearly if S-TAGGED frames are to be admitted at all then the priority needs to be under the control of the provider. The only other option is to "Admit Only Untagged". This loses some value in some interfacing scenarios as it means that equipment in the provider's physical domain has to get involved in flexibly mapping C-TAGs to S-TAGs, whereas in some scenarios it would be rather more convenient for the customer to both carry that burden and be able to change the mapping without talking to the service provider. Perhaps the customer has their own Provider Edge Bridge to perform the mapping (don't confuse ownership and functionality). What is more there may be customer scenarios, particularly with high end customers using routers to attach to their customers, where introduction of the C-TAG would be a completely artificial step. It seems then that the S-TAG and the priority both need to be policed, and the rule we have to change is to allow remapping of priority, at least at Provider Bridge Ports that are Customer Network Ports (ports labelled 2 and 3 in the proposed disposition to comment 5 on port naming). It seems likely that regeneration on already tagged frames will be required for proper support of drop precedence etc. etc. Further a service provider may reasonably want to provide instances of service to a customer that wishes to use each of those instances to support distinct C-VID spaces, without requiring each space to be a different subset of the whole. The finished standard should not artificially limit the services to be provided, beyond ensuring correct service selection. segregation, and identification.

The architectural argument that one solution is better than the other, by virtue of one ensuring that C-TAGs are only understood by C-VLAN aware Bridges and S-TAGs only by S-VLAN aware Bridges is specious. Either the C-Bridge has to add an S-TAG for the S-Bridge to pickup the priority, or the S-Bridge has to scan for and interpret a C-TAG. In both cases the TAG is not being understood by the Bridge for which it is not native in the sense that a VLAN-ID is neither being inserted or removed from the TAG, the only item of interest being the priority (which can be passed through the ISS - no new interfaces!). The situation is marginally worse when the S-Bridge scans for the C-TAG as it actually has to recognize and check the TAG, while when the C-Bridge adds the priority it just has to insert one of eight fixed 32-bit patterns into the frame. However it is clear that in an environment of existing customer equipment, the use of C-priority by S-Bridges has more application.

The C-Bridge is essentially functioning as an end-station from the point of view of the S-Bridge, so the argument about a bridge not transmitting priority-tagged frames has little force, particularly as the C-Bridge doesn't think the priority pattern added into the frame is a TAG.

The result of this discussion is that we should:

(a) specify optional priority extraction from a C-TAG (as Steve suggests in this comment), for use as an ISS component in the MAC stack of an S-Bridge. This should be a new clause 6.x, after the current clause 6.8.

(b) should recognize and document that customer systems can send S-TAGged frames to convey priority to a provider S-Bridge

(c) explicit permit user priority regeneration on receipt in an S-Bridge. The decision as to how exactly this feature is to be specified and whether it equally applies, with some caveats, to a C-Bridge will be affect by the result of the discussion on drop precedence

in addition:

(d) we may or may not decide to support option (b) by explicitly retaining the current priority clause 6.8. I suggest that we do so, though the clause heading needs changing.

(e) we have a choice as to whether in specifying (a) we actually require each frame to be priority tagged, i.e. not carry a VID (we remove the tag and discard the VID anyway). I suggest that we take a permissive approach, i.e. allow the VID rather than a null VID to be present. I think this opens up few possible mistakes for the customer.

Comment 13 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: ER
CLAUSE: 8.6
PAGE: 41 LINE:
COMMENT START:
It is unclear how and when priority regeneration occurs and whether it should be part of the VLAN translation table as well (i.e. is there a priority translation table other than the priority regeneration table)? In 802.1D priority regeneration is talked about prior to talking about the forwarding process. In this document many of the ingress activities are discussed in the forwarding process clause. I would assume the priority regeneration occurs in the ISS as before, but since VLAN translation occurs later in a Provider bridge, should a separate priority mapping be preformed as well?
COMMENT END:
SUGGESTED CHANGES START:
Include user priority regeneration in the discussion of 8.6, and if appropriate include it in the figure 8-8 and discuss. Consider whether the VLAN translation table should also translate priority.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 13

Reject the specific suggestion - but see below for improvements to address the clarity issue. See also the proposed resolution to comment 12.

Priority regeneration is carried out by functions in the receiving port's "MAC stack", not in the relay entity. Any priority regeneration is already done before the data indication (DA, SA, priority, VID, data) gets into the MAC Relay Entity - the relay entity deals with the parameters of the indication, not directly with their encoding in a frame format.

This means the suggestion to discuss priority regeneration in 8.6 should be rejected, but the priority regeneration for media access method independent support of the ISS should be discussed in the relevant clause 6.x. Clause 6.8 in D2 is, for example, deficient as it just describes adding priority on transmission and not the reception part of the processing - which is to be found in the support of the EISS - so clause 6.8 should mention that. If a new Clause 6.9 were to be added, as suggested for comment 12, then recovery of the priority from the C-TAG as part of the S-Bridge "MAC stack" would do the priority regeneration below the ISS. Note in passing that this formulation also defends the provider network against a malicious customer sending a frame that begins C-TAG.S-TAG.data in an attempt to provide an S-TAG that won't be checked by the provider, and thus gain unauthorized access to a service instance. Conversely if the C-priority tagged frame was simply identified as untagged and passed right the way to clause 8.6 an attack could be facilitated (a very strict implementation would hold the S-VID separate from the data and not be attackable, but most would not and thus admit the attack).

Comment 14 Arjan De Heer

NAME: Arjan de Heer
COMMENT TYPE: T
CLAUSE: 5.6.2 and 6.8
PAGE: 12 LINE:
COMMENT START:
Clause 5.6.2. introduces the option that a customer bridge does priority selection as specified in 6.8. This section describes that the customer bridge can signal the priority to the provider, by adding an S-tag to the frame (if I understand 6.8 correctly). In other words the translation of customer priority into service provider priority is done at the customer equipment. The same functionality can be realized at the provider side. The provider maps the C-tag priority into an S-tag priority. This also allows existing vlan aware bridges to signal the priority to the service provider.
COMMENT END:
SUGGESTED CHANGES START:
Remove this functionality from the customer bridge, put into provider bridge
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 14

See proposed disposition of comment 12.

3.5 BPDU handling by C-VLAN aware component

Comment 15 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: ? PAGE: ? LINE:
COMMENT START:
Somewhere, I'm not sure if it's in requirements in Clause 5, or in descriptions of Provider Bridge operation in clause 15 or 16, there needs to be more discussion of how Customer BPDUs are handled by Provider Bridges. I think it is clear that when a Customer BPDU is received at the Customer Network Port of a Service-VLAN aware Bridge component, it is tagged with a S-TAG and forwarded like any other customer frame. I can think of two specific cases, however, that need to be more explicit. The first is how Customer BPDUs are handled by the Customer-VLAN aware Bridge component of a Provider Edge Bridge. One option is that this component processes Customer BPDUs as a peer in the customer network, and generates Customer BPDUs to be transmitted through each Internal Port to a Customer Network Port of the Service-VLAN aware Bridge component, where they are handled as described above. Alternatively there may be a way for the C-VLAN aware component to replicate the BPDUs to each Internal Port without actually processing them (?). Also there has been a proposal that the Customer-VLAN aware Bridge component not process BPDUs but forward them to the single Internal Port that would be used for untagged or priority tagged frames. This has significant ramifications on how the provider network affects customer spanning trees. The second case is where a customer packets received at the provider network contain Service Tags (from a customer operated Provider Bridge). Must all packets received be tagged with an S-Tag? What happens to non-S-Tagged packets? Specifically how should untagged customer BPDUs handled in this case?
COMMENT END:
SUGGESTED CHANGES START:
Discuss. My preferences are: In the first case, Customer BPDUs are processed by the Customer-VLAN aware Bridge component of a Provider Edge Bridge. In the second case, all packets received at a Provider Bridge from a customer operated Provider Bridge must contain S-Tags. Untagged packets are discarded. Customer BPDUs must contain an S-Tag that signifies the service intended to carry that packet.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 15

Accept - broadly but with one exception. To treat processing of C-BPDUs by a C-Bridge differently just because that C-Bridge belongs to a provider is just to reopen the possibility of the control plane protocols being treated incorrectly. The discussion to reinforce this probably belongs in clause 16.

The exception is that a Provider Bridge providing a single service instance on a single physical port does not and should not need the packets to contain an S-TAG. All frame are allocated to the service instance and tagged. The C-BPDUs, though they may lack an S-TAG carry a different multicast address than S-BPDUs so they will simply be S-tagged by the Provider Bridge, with the tag for that service instance. This works just fine.

3.6 Conformance

Comment 16 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 5.8.2
PAGE: 13
LINE:
COMMENT START:
Why are we talking about the conformant of a component in here instead of conformant of a system. The conformant for Service-VLAN aware Bridge component was defined in the clause 5.7.2.
COMMENT END:
SUGGESTED CHANGES START:
SUGGESTED CHANGES END

Editor's Proposed Disposition of Comment 16

Discuss. I don't understand the comment.

3.7 Provider Bridge Description

Comment 17 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 5.9
PAGE: 13
LINE:
COMMENT START:
I think this is a very nice attempt in describing the Provider Bridge in terms of its components. However, for better clarity it would be good to move up the picture of the clause 16.9 into this clause and show different port types (e.g., Provider Edge port, Customer Network port, and Provider Network port) on the diagram - to show the port types for CPE1, CPE2, and CPE3 of the figure 16-1. It would also help to explain that if a Provider Bridge comprises of only Service-VLAN aware bridge component, then there is no Provider Edge port.
COMMENT END:
SUGGESTED CHANGES START
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 17

Copy some of the (elements of the) simple pictures in clause 15, i.e. figures 15-3, 15-4, into Clause 5 so that the user does not have to redraw these from the text. The pictures in clause 16 are really intended to describe how a network can be built from the components, and should remain there.

3.8 Drop precedence

Comment 18 Dinesh Mohan

NAME: Dinesh Mohan
COMMENT TYPE: TR
CLAUSE: Editors' Foreword
PAGE: b
LINE: 1
COMMENT START:
Though it has been noted by the editor that drop precedence related text has not been incorporated, it was generally agreed that both options i.e. a) use of p-bits to also indicate drop precedence, and b) use of CFI to indicate drop precedence will be included. It is recommended that the discussion be reflected in this document to facilitate further discussion and resolution.
COMMENT END:
SUGGESTED CHANGES START:
Include as comments the two proposals on table for drop precedence as were discussed during last meeting.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 18

As discussed at the January 2004 interim.

It was definitely stated at the beginning of the work that completely rearchitecting a bridge was out of bounds - no new major building blocks, no new service interfaces. As such the editor/task group chair requires a substantive complete proposal that the whole group can formally approve before risking getting into a tar pit (it was previously thought that proposals could be expanded without the tar pit, but that proved not possible). Accordingly the editor asked a small subgroup, including those proposing the two drop precedence alternatives, to work out the extent of the impact and ramifications of drop precedence throughout the document, instead of putting in a large volume of change and then taking it out again.

Comment 19 Mohan, Parsons, Bottorff (Drop precedence - service parameter)

NAME: Dinesh Mohan
COMMENT TYPE: TR
CLAUSE: 6.6.1
PAGE: 22 LINE:
COMMENT START:
The EISS should support additional parameters. These include: drop precedence, both customer and service provider's FCS.
COMMENT END:
SUGGESTED CHANGES START:
Add parameters:
- drop_precedence
- s_frame_check_sequence
Change, existing parameter:
- frame_check_sequence to c_frame_check_sequence
SUGGESTED CHANGES END:

NAME: Glenn Parsons

COMMENT TYPE: TR (Technical, Required)
CLAUSE: 6.6.1
PAGE: 22 LINE:
COMMENT START:
Drop precendence needs to be added to the draft as noted in the editor's notes. The EISS service primitives is one such place.
COMMENT END:
SUGGESTED CHANGES START:
Add a new parameter 'drop_precedence' to both EM_UNITDATA.indication and EM_UNITDATA.request and add the following descriptive text:
The drop_precedence parameter carries the drop precedence.
SUGGESTED CHANGES END:

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 9.7
PAGE: 76 LINE:
COMMENT START:
Drop precendence needs to be added to the draft as noted in the editor's notes. The S-TAG is one such place.
COMMENT END:
SUGGESTED CHANGES START:
Add a single bit 'drop_precedence' the S-TAG.
SUGGESTED CHANGES END:

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 6.6.1 PAGE: 22
LINE: EM_UNITDATA.indication
COMMENT START: The EISS should support Drop Precedence. The EISS needs a parameter to carry the Drop Precedence added to the EM_UNITDATA.indication and EM_UNITDATA.request. COMMENT END: SUGGESTED CHANGES START: Add a parameter called drop_precedence to the parameters of EM_UNITDATA.indication and EM_UNITDATA.request. In addition add to 6.6.1 a definition that the drop_precedence parameter carries the drop precedence.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 19

See proposed disposition to comment 18.

3.9 Additional FCS

Comment 20 Parsons, Bottorff, Chen

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 6.7.2
PAGE: 25 LINE:
COMMENT START:
This standard must support the ability to retain customer FCS at the egress of the provider network.
This requires several additions:
- addition of a new provider FCS parameter
- description of the semantics
COMMENT END:
SUGGESTED CHANGES START:
Modify text as follows:
The value of the frame_check_sequence parameter is determined as follows:

f) If the service_frame_check_sequence parameter is present, the frame_check_sequence parameter value is used as-is. Otherwise;
g) If the frame_check_sequence parameter received in the data request is either unspecified or still
carries a valid value, then that value is used. Otherwise;
h) The value used is either derived from the received FCS information by modification to take account
of the conditions that have caused it to become invalid, or the unspecified value is used.
The value of the service_frame_check_sequence parameter is determined as follows:
i) If the service_frame_check_sequence parameter received in the data request is either unspecified or still
carries a valid value, then that value is used. Otherwise;
j) The value used is either derived from the received service FCS information by modification to take account
of the conditions that have caused it to become invalid, or the unspecified value is used.
NOTE-The original FCS associated with a frame is invalidated if there are changes to any fields of the frame, if fields are added or removed, or if bit ordering or other aspects of the frame encoding have changed. An invalid FCS is signalled in the E-ISS by an unspecified value in the frame_check_sequence or service_frame_check_sequence parameter of the data request primitive. This signals the need for the corresponding FCS to be regenerated according to the normal procedures for the transmitting MAC.
The options for regenerating the FCS under these circumstances are discussed in Annex G of IEEE Std 802.1D, 1998 Edition.
SUGGESTED CHANGES END:

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 6.6.1
PAGE: 22 LINE: EM_UNITDATA.indication
COMMENT START:
The EISS should support a service frame check sequence(S_FCS) which is independent and in addition to the customer FCS(C_FCS). Provider Bridges should use the S_FCS for checking provider bridge frames and should not modify the C_FCS.
COMMENT END:
SUGGESTED CHANGES START:
Change the name of the frame_check_sequence parameter of the EISS EM_UNITDATA.indication and EM_UNITDATA.request to C_frame_check_sequence. Add a parameter to the EM_UNITDATA.indication and EM_UNITDATA.request called S_frame_check_sequence. Add a description of the S_frame_check_sequence which states that the S_frame_check_sequence carries the FCS of a provider frame. The S_frame_check_sequence is generated whenever a 802.3ad bridge receive a customer frame. When using the S_frame_check_sequence the C_frame_check_sequence is placed at the end of the data field and is never modified by the provider bridge.
SUGGESTED CHANGES END:

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 8.6.1
PAGE: 42 LINE:
COMMENT START:
Provider bridges must retain the customer frame check sequence unmodified. This section should describe how a new Service Frame Check Sequence is created for a customer frame.
COMMENT END:
SUGGESTED CHANGES START:
Add text after paragraph 5 indicating that when a provider bridge receives a frame with a s_frame_check_sequence of NONE, then the provider bridge will generate a s_frame_check_sequence.
SUGGESTED CHANGES END:

NAME: Michael Chen
COMMENT TYPE: TR
CLAUSE: All
PAGE: All
LINE: All
COMMENT START:
Current proposal does not provide the option to carry customer Ethernet FCS. This is a contentious topics as indicated by the mailing lists. However, multiple service providers with deployment experience have voice strong support in adding an option to carry customer FCS. Since the target users of this technology are the service providers, 802.1ad should look to provide a mechanism to optionally support this behavior.
COMMENT END:
SUGGESTED CHANGES START:
Propose to have 32 bit S-Tag with 24 bits assigned to be S-VID. Two bits will be used for payload indication including service OAM and optional customer FCS retention. The rest 6 bits will be reserved. The reserved 6 bits may be used in the future for things such as drop precedence and QoS indication.
By having a field in the expanded S-Tag indicating presence of customer FCS retention, we satisfy the needs voiced by the service providers. Furthermore, we also provide the capability for future expansion.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 20

Discuss. This proposal (FCS) is outside the bounds stated at the beginning of the project at PAR discussion time (no new major blocks, no new service interface parameters). Such changes would likely very considerably extend the lifetime of the current project. The proposed change would have radical effects on existing equipment operation, interoperability, and coincidentally on frame size. Discussion on the email reflector did not seem, on balance, to support the FCS addition, and noted that such a change would require quite a different philosophy towards the internal operation of bridges. With something of this major impact and change a substantive proposal backed by a formal WG decision is appropriate, and as stated at the first session of the January 2004 meeting, a proposal to the group and a formal vote to extend the effective project is appropriate. If there is support for starting work on such a change, a vote could be held during the closing 802.1 plenary (Thursday) in March. This would need to be followed by significant liaison with 802.3 as a change in the way FCS's are used can be expected to cause them considerable concern. Note that the deployment experience cited relates to the architectural models used to support circuit services. The discussion in the IETF PWE3 group that set this all off was started by a single bridge failing in a lab, and the anticipated consequences if that should be deployed in a network.

3.10 Acceptable frame types

Comment 21 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: 8.6.1
PAGE: 42 LINE:
COMMENT START:
I believe we will need a "Admit only un-tagged frames" value for the Acceptable Frame Types parameter. Specifically this will be the value used for Customer Network Ports of Service-VLAN aware Bridge components (since these are service VLAN aware, the un-tagged means the frames do not contain a service tag, but may or may not contain a customer tag).
COMMENT END:
SUGGESTED CHANGES START:
Add "c) Admit Only un-tagged frames". This value may be optionally supported by Customer Bridges. Discuss whether it is optional or required for Provider Bridges.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 21

Accept/Discuss. This change significantly reduces the configuration requirement in certain simple scenarios. Suggest optional support by Customer Bridges, required by Provider Bridges.

3.11 Reserved Addresses

Comment 22 Dirceu Cavendish

NAME: Dirceu Cavendish
COMMENT TYPE: TR
CLAUSE: 8.6.4
PAGE: 44 LINE:
COMMENT START:
Shouldn't the reserved addresses of Tables 8-1 and 8-2 be different?
COMMENT END:
SUGGESTED CHANGES START:
Correct addresses.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 22

One table is a subset of the other. This is particularly what is required to support the layer model.

3.12 Service instances can translate between interface types and selectors

Comment 23 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: 15.1
PAGE: 84 LINE:
COMMENT START:
The service provided by Provider Bridges to attached Customer Bridges is not always transparent to the operation of media access method independent functions by the Customer Bridges. Specifically, the C-TAG transmitted to a destination customer system from a Provider Edge Bridge may or may not be the same as the C-TAG received by a Provider Edge Bridge from the originating customer system.
COMMENT END:
SUGGESTED CHANGES START:
Change the second paragraph of 15.1 to: "The service provide by Provider Bridges to attached Customer Bridges is transparent to the use of the MAC Service provided to end stations on LANs attached to the Customer Bridged LANs. The service is transparent to the operation of media access method independent function by the Customer Bridges, with the exception that the Customer VLAN Identifier may be modified by the service in cases where the C-VID is used as a service instance selector (15.4)."
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 23

Accept in principle, replacing "used as a service instance selector (15.4)" with "used for service instance selection and identification" adding appropriate references, and further modifying the text so that it is clear that the customer is aware of the change and the change happens on all frames. There may be better ways of phrasing the same intent. The intent of the comment appears clear.

3.13 Provider tagged service interface

Comment 24 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: 15.5
PAGE: 87 LINE:
COMMENT START:
It is not immediately clear what is intended to be highlighted by Figures 15-6 and 15-7.
COMMENT END:
SUGGESTED CHANGES START:
In Figures 15-6 and 15-7 change the label of the right most bridge to "Customer Operated Provider Bridge".
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 24

Accept.

3.14 Service selection

Comment 25 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: T
CLAUSE: 16.3
PAGE: 91 LINE:
COMMENT START:
The bulk of this clause talks about service selection and identification by a Provider Bridge based on the VID received in an S-TAG. In Note 9 and the previous paragraph it starts talking about a customer system selecting and identifying a service. The statements are all true and consistent for the case where the customer system is a customer operated Provider Bridge (a la 15.5), but does not differentiate (and the statements are inconsistent with) the case where the customer system is a Customer Bridge using the VID of a C-TAG for service selection (15.4), or Port-based service selection (15.3).
COMMENT END:
SUGGESTED CHANGES START:
Clarify that the VID translation table is used only for the VID of S-TAGs. Distinguish between customer selection using S-VID versus C-VID.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 25

Discuss. This clause, and NOTE 9 in particular needs substantial rework to come up to date with the "2-in-1" Provider Edge Bridge model. I think that VID translation could be permitted for a VID of any type (if the VIDs are of different sizes a VID extension rule for the smaller would need to be adopted), but that should come out after Clause 16 is brought up to date.

Comment 26 Dirceu Cavendish

NAME: Dirceu Cavendish
COMMENT TYPE: TR
CLAUSE: 15.4
PAGE: 86
LINE: Two last paragraphs
COMMENT START:

Explanatory text of Fig. 15-4 does not agree with figure's content.
COMMENT END:
SUGGESTED CHANGES START:
Replace last paragraphs with:
"The customer-VLAN aware Bridge supporting the customer-tagged interface assigns each frame received from the attached customer system to a Customer VLAN, and presents it at each port-based interface that supports a service instance that the customer desires to carry that VLAN. Similarly, frames received from the S-VLAN aware Bridge are assigned to a Customer VLAN and hence to one or more customer-tagged interfaces.
The Customer-VLAN aware Bridge provides service instances for a single customer. The S-VLAN aware Bridge can provide port-based interfaces to a number of Customer-VLAN aware Bridges serving different customers."
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 26

Accept.

3.15 VID strong typing

Comment 27 Glenn Parsons

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 6.6.1
PAGE: 23
LINE:
COMMENT START:
Indicating that the vlan_identifier simply carries VID is too vague given that in can be a customer or service VID.
COMMENT END:
SUGGESTED CHANGES START:
Reword as follows:
The vlan_identifier parameter carries the customer VLAN identifier (C-VID) or service VLAN identifier (S-VID) depending on the particular instance per 6.7.
SUGGESTED CHANGES END:

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 6.6.1
PAGE: 23 LINE: 1
COMMENT START:
The definition of the vlan_identifier parameter needs to indicate that it may carry either the C-VID S-VID.
COMMENT END:
SUGGESTED CHANGES START:
Change definition of vlan_identifier parameter to: The vlan_identifier parameter carries either the C-VID or the S-VID depending on the instance of the EISS.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 27

Reject. I don't see why this parameter needs to be so strongly typed. A VID is practically an integer. The proposed changes simply add part of a description that can happily be elsewhere in the document and place it at this particular point (6.6.1). If this practice, of following ramifications and explaining alternative, is followed consistently any document rapidly becomes unmanageable. The proposed rejection is quite independent of whether C-VIDs or S-VIDs have the same size.

3.16 Specify protocol use of frame formats in frame format section

Comment 28 Paul Bottorff

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 6.7
PAGE: 23
COMMENT START:
The customer and service frame formats are different, the second sentence of 6.7 does not state that the customer and service frame formats are different and that EISS instantiation shall support one frame format of the other but not both.
COMMENT END: SUGGESTED
CHANGES START:
Change the second sentence of 6.7 to read: Any instance of the EISS shall be supported by using one but not both of the customer frame format or the service frame format. In addition, any given instance of the EISS shall support one but not both of the following VLAN tag types: a) Customer VLAN tag (C-TAG) for customer frame format; or b) Service VLAN tag (S-TAG) for service frame format.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 28

Reject. The initial paragraph very clearly states one but not both of the tags. What is at issue here is parameter mappings and usage, not formats. Formats are specified in Clause 9. There is no need to replicate Clause 9 information here.

3.17 S-VID size

Comment 29 Muneyoshi Suzuki

NAME: Muneyoshi Suzuki
COMMENT TYPE: TR
CLAUSE: 9.7
PAGE: 76 LINE: 5
COMMENT START:
The following comment is a repeat of comment 44 in the disposition of ballot comments on P802.1ad/D1. This is because, as discussed on the mailing list, this comment does not require globally unique VID space. However, this comment has not discussed P802.1ad ballot resolution session in the meetings, so the misunderstanding is not straighten out. 12bit VID space proposed for Service VLAN Tag Control Information is too small for support of public service, so extension is needed. It should be unique in a single Provider Bridged Network, but does not require globally unique VID space.
COMMENT END:
SUGGESTED CHANGES START:
Extend VID to 24 bit. If this is not desired, clarify technical reasons.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 29

Discuss. See also proposed resolution of comment 30.

The need for single service provider to be able to offer more than 4k service instances is not in doubt. The significant questions include

(a) how many Provider Bridge Networks make up the service provider's network

(b) are the service offerings for high end business customers for their site to site connectivity intended to be the same as for residential or low end business connectivity to a local internet provider. (For internet access connectivity one C-VID per customer inside an S-VID supporting 4k customers is enough, thus supporting up to 16 million internet connections).

There is probably a considerable difference in philosophical approach between those who wish to put together a structured network to take advantage of the VLAN bridging technology internally, and those who don't want to be bothered by those details. For those taking a structure approach some detail on the expected distribution of customers and connectivity across the resulting service provider network would be highly useful in reaching a conclusion as to whether 12 bit VIDs with local renumbering would work or not. There are significant anticipated gains from being able to interwork with existing 12 bit VID technology, one intermediate position might be to allow Provider Bridges the option of using a larger VID space internally - thus truly supporting 4K VLANs per port. With the current limit on MSTP spanning trees of 64 (see also comment 30) that can be easily used to provide up to 256K VLANs before spatial reuse of VIDs on each tree. This is probably the maximum number of VLANs that any one would wish to see passing through a single item of equipment.

Some data on how many VLANs would be required to be supported by any single item of equipment would be very useful to take this discussion forward. I presume everyone thinks that 2**24 is unrealistic, so that the equipment controlling the network will have to have traffic engineering components to associate each service instance with a particular network path (or set of potential paths). Some participants clearly believe that that equipment can also handle local VLAN numbering/circuit assignment within a structured network, there may be data they do not have as to why that is not feasible. We have had quite a few presentations on the use of provider bridging islands within a larger network, so far we haven't had any formal presentations on the assumptions and use of a very large flat network taking into account the limitations of equipment. Additionally if there is further information that those who have presented on the "islands" approach could provide to the rest of the group then it would be useful to have that specifically requested as well.

Comment 30 Parsons, Bottorff, Mohan, Chen

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 9.7
PAGE: 76 LINE:
COMMENT START:
The figure is a copy of Figure 9-1 and should be Figure 2 S-TAG TCI format. Further, given the larger S-VID needed, the S-TAG should be increased to 32 bits.
COMMENT END:
SUGGESTED CHANGES START:
Rename figure "Figure 2 S-TAG TCI format" and redraw figure to show 24 bit S-VID, 3 bit priority, 1 bit drop precedence and 4 reserved bits.
SUGGESTED CHANGES END:

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 9.7
PAGE: 76 LINE:
COMMENT START:
It is expected that service providers will want to build 'islands' that can support more than 4000 customers, each with that many service instances. The result is that current VID tag is inadequate to address this requirement. To allow for some growth it is proposed to make the S-VID 24 bits
COMMENT END:
SUGGESTED CHANGES START:
Show the S-VID as 24 bits in length.
SUGGESTED CHANGES END:

NAME: Dinesh Mohan
COMMENT TYPE: TR
CLAUSE: 9.7
PAGE: 76
LINE:
COMMENT START:
The proposed semantics and structure of Service VLAN TAG are claimed to be identical to Customer VLAN TAG with the exception of bit 5. However, this imposes a limitation of 12-bit service instances which conceivably are sufficient for customers but is limiting for service instances that can be offered by service providers who would make use of Provider bridge specification. It is recommended that S-TAG structure should be changed such that number of service instances can be increased. It is suggested that Service VLAN TAG structure should be such that 24-bits are allocated for S-VID.
COMMENT END:
SUGGESTED CHANGES START:
Change the S-TAG TCI format in figure 9-1 (change figure text to "figure 9-2") to a 32 bit structure with 24 bit S-VID, 3 bit service priority, 1 bit drop precedence, and 4 bit reserved fields.
SUGGESTED CHANGES END:

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 9.7 PAGE: 76 LINE:
COMMENT START:
Add a drop precedence indicator to the S-TAG.
COMMENT END:
SUGGESTED CHANGES START:
Reserve an additional bit for DP in a 32 bit S-TAG.
SUGGESTED CHANGES END:

NAME: Michael Chen
COMMENT TYPE: TR
CLAUSE: All
PAGE: All
LINE: All
COMMENT START:
Current proposal of 12 bit S-VID only support a limited number of service instances (4K). This number is simply too small and will not be able to support service providers' needs. With the current 12 bit S-VID field proposal, there exist a high probability a customer tagged service interface consuming all the service instance in a provider network due to the same S-VID and Customer VID space.
The use of multiple provider bridge "islands" does not really address this issue and further introduces architecture and operational complexity. How do one connect provider bridge "islands" together gracefully? By having a larger S-VID, this issue is largely removed.
COMMENT END:
SUGGESTED CHANGES START:
Propose to have 32 bit S-Tag with 24 bits assigned to be S-VID. Two bits will be used for payload indication including service OAM and optional customer FCS retention. The rest 6 bits will be reserved. The reserved 6 bits may be used in the future for things such as drop precedence and QoS indication.
By having a larger S-Tag and S-VID, we simplify the provider bridge operation. Furthermore, we also provide the capability for future expansion.
SUGGESTED CHANGES END:

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 9.7
PAGE: 76 LINE:
COMMENT START:
The 4K MAC service instances supported by the 12 bit S-TAG is too few.

COMMENT END:
SUGGESTED CHANGES START:
Change the S-TAG in figure 9-1 to a 32 bit field. Allocate 24 bits for the S-VID.

Editor's Proposed Disposition of Comment 30

Discuss.

At the beginning of the provider bridging work, some participants contended that an 802.1D bridging like model was completely inappropriate for the scaling requirements of a provider bridged network, and that an encapsulation scheme (MAC-in-MAC) would have to be used right across the provider network as opposed to a doubly tagged (Q-in-Q) scheme. Since encapsulation always introduces a whole set of new problems (thinking back to the now obsolete FDDI encapsulating bridges) it is pleasing to note the new enthusiasm of those participants for very large, flat, transparent network supporting up to 16 million service instances, and their rejection of network structure (including encapsulation schemes).

Hopefully the change of heart has something to do with the presentation and paper showing that address learning is scalable, and that the total address learning is a function of the customer attachments and is independent of the organization/internal size of the provider bridged network.

However those who advocated the use of Q-in-Q initially appear to be far more conservative in their deployment of this technology. The initial thrust, of at least some of the arguments, was to use this very cost effective technology towards the edge of the network, where the number of systems is large, and facilitate scaling to the next order of magnitude with relatively few MPLS core systems (with as few as one four thousandth of the number of MPLS ports required, by virtue of the Q-in-Q layer). Of course the service perceived by the attached customers is exactly the same whether there is or is not such an MPLS core, or whether the core functionality is provided by something else.

This latter, structured, approach to network construction gains most of the benefit of using VLAN Bridges to support such a service without having to radically scale the control plane specified by 802.1. For example it is not necessary to work out how to get GVRP to perform registrations for 16 Million service instances. On the first ballot we accepted a comment to the effect that 64 spanning trees would be enough for .1ad, as performing traffic engineering by configuring spanning trees was of limited interest for very large networks - which would rather use traffic engineering tools built for MPLS (or similar). If we are to handle 2**24 service instances we will have to scale MSTP to handle much more in the way of traffic engineering, in an environment where transmission facilities cost at least an order of magnitude more than the original design environment.

It would thus seem that simply adopting a proposal for a large S-VID size without any accompanying proposal for scaling control elements (or indeed without any survey of how those control elements ought to be scaled) would put us in the situation of having suddenly acquired a large amount of new work with no one willing to do it (those who believe in layering provider bridges around a core of a different type being unlikely to volunteer for that work). This would then risk indefinitely postponing the provider bridging work as it would then once more be attacked as not scaling. I am sure that this is not the intent of the comments.

A survey of the control plane scaling challenges by those who would like a 2**24 S-VID, together with detailed proposals for protocol enhancements in each area, would be most helpful before we accidentally find ourselves in the above trap.

Data on the real scaling limits of a network that locally uses 12-bit VIDs/Connection numbers, with number reuse allowed, would also be useful.

3.18 Instances of the MAC Service

Comment 31 Muneyoshi Suzuki

NAME: Muneyoshi Suzuki
COMMENT TYPE: TR
CLAUSE: 8.2
PAGE: 35 LINE: Figure 8-2
COMMENT START:
VLAN-aware Bridge architecture shown in Figure 8-2 is inconsistent with Bridge architectures defined in 802.1D-1998, 802.1D/D4, and 802.1Q-2003. The problem is that the MAC Service is provided by the Internal Sublayer Service in the Figure. However, the current 802.1D and 802.Q specify that the MAC entity directly supports both MAC service and ISS. The first line of Clause 6.4 in 802.1D-1998 explicitly addresses that the ISS support the MAC Relay entity. It is also described in the same Clause of 802.1D/D4 and it is not amended in 802.1Q-2003. And Clause 6.5 in 802.1D-1998 defines support of the ISS by specific MAC procedures defined in the series of 802 standards. Furthermore, ISO/IEC 15802-1 (1995) defines the generic MAC service and mapping to/from specific MAC procedures. However, neither 802.1D nor 802.1Q defines a procedure that maps the generic MAC service to/from the ISS. And it is useless to provide generic MAC service supported by ISS implemented with specific MAC procedure. Therefore as shown in Figure 7-3 in 802.D-1998 and 802.1D/D4 and Figure 8-3 in 802.1Q-2003, the MAC entity directly provides both MAC service and ISS.
COMMENT END:
SUGGESTED CHANGES START:
Fig. 1 in the attached ppt file illustrates proposed VLAN-aware Bridge architecture based on the definitions of current 802.1D-1998, 802.1D/D4, and 802.1Q-2003. MAC Entity directly supports both MAC service and ISS. And E-ISS is provided by ISS. However, based on the sublayer architecture proposed in Clause 6 of 802.1ad/D2.0, the MAC Entity should be separated sublayer functions. Therefore, proposed VLAN-aware Bridge architecture based on sublayer functions is shown in Fig. 2. A MAC Entity consists of MAC Specific Function (MSF), MAC Dependent Convergence Function (MCF), and MAC Independent Function (MIF). It is obvious from the first line of Clause 6.5 in 802.1D-1998 and 802.1D/D4 that the MCF supports ISS but does not directly support MAC service.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 31

As discussed at the January 2004 meeting, and in part before the email (this proposed disposition summarizes for the record). See also the proposed disposition of comment 32.

P802.1ad is an amendment to 802.1Q, so it can change what 802.1Q says - as required of course by the Provider Bridges effort (not just because we have a good idea).

The problem is that the current 802.1D specification is not specific enough, and in just stating that the specific MAC supports both service interfaces, does not specify at all how that it is done. In particular the current specification provides no clue as to how an ISS request can get turned into an MS indication if the request has a destination address that matches an address that the LLC client of the MS wants to receive. Now that can be spelled out (in clause 8.5). Thus the problem with the current 802.1D specification is that it is not able to operate correctly with "shims" (new media independent functions) inserted into the stack. Not only does this problem occur with P802.1ad, it also happens with MACsec (P802.1AE). With MACsec in particular it is impossible to cheat as 802.1D does, since security is required for the MAC Service.

While support of all the end station functions provided within a bridge are clearly within the scope of the 802.1Q specification, a thorough fix would specify the support of MAC Service in end stations, including use of VLANs within end stations. This would be better done, as discussed in the meeting, as part of the P802.1ac MAC Service revision work, rather than putting a large (and fairly obvious) aside in .1Q. It has become clear that we will need to take rather more care about the "MAC stack" in Bridge Ports (an end

station with MAC Security). In the past we have gone along with, to avoid fruitless argument, a number of cases where a protocol (like a media access control) is confused with a service realized using some trivial procedures plus the protocol. Attached to a bridge port there can be several higher layer entities, using different SAPs with some (like LLDP) not using a VLAN shim, while others (like an IP stack supporting SNMP) do. There is no conceptual problem supporting several different MSAPs with these different characteristics over a LAN, and their support by a single bridge port behaves the same. Over time the MAC Service (15802-1 and other versions) and the ISS have converged somewhat (P802.1ad helps this convergence as some of the MACs that were decidedly "odd' have been retired.)

Look forward to more diagrams, especially as we pursue Connectivity Fault Management, showing the precise relationship of the higher layer entities in a bridge to the underlying service.

Comment 32 Muneyoshi Suzuki

NAME: Muneyoshi Suzuki
COMMENT TYPE: TR
CLAUSE: 8.2
PAGE: 35 LINE: 4-6
COMMENT START:
I'm not sure semantics of "Bridge Port providing multiple instances of the MAC Service". According to ISO/IEC 15802-1 (1995), a set of MAC Service primitives are issued at an MSAP, and an MSAP is identified by a MAC Address. And Clauses 7.2.3 of 802.1D-1998, 7.2 of 802.1D/D4, and 8.2.3 of 802.1Q-2003 specify a MAC entity corresponds to a Port. Therefore, "Bridge Port providing multiple instances of the MAC Service" implies that the single MAC entity has multiple MAC addresses, but I'm not sure such MAC technology.
COMMENT END:
SUGGESTED CHANGES START:
There may be misunderstanding, so clarifications are needed.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 32

See the disposition of comment 31 for a plan of action. From an architectural modelling point of view a MAC that is capable of promiscuous receive and of transmitting with any source address (as any MAC that supports the ISS to support bridging must be) can support a number of MSAPs identified by their own address/VLAN. From a conservation of addresses (registration authority concern) the number of addresses is kept as close to 1 as possible. Implementations of MACs solely for end station use may only be capable of supporting one address, but that's of no immediate concern in this standard.

3.19 Addressing of higher layer entities

Comment 33 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: E
CLAUSE: 8.14.9
PAGE: 68 LINE:
COMMENT START:
In figure 8-11, two independent higher layer entities are shown attached to two different ports and not using the relay to allow communication of those higher layer entities to other ports in the system. It would be valuable to remind people that in this particular case, it is important that the higher layer entities are using the source MAC address of the port they are attached to. The new management port concept allows us to have some higher layer entities use a common source MAC address, but it must use the relay function for this to work.
COMMENT END:
SUGGESTED CHANGES START:

Add a sentence or note to the last sentence that indicates entity A and entity B are using the source MAC address of the port they are communicating out.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 33

Accept.

3.20 Provider Bridge positioning within MAC Sublayer

Comment 34 Glenn Parsons

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 15
PAGE: 84 LINE:
COMMENT START:
Figure 15-1 is not entirely clear or accurate especially on the bookends. Notably the 'imaginary MAC' concept that had been agreed is missing.
COMMENT END:
SUGGESTED CHANGES START:
Replace Figure 15-1 with the slide presented by Paul Bottorff at the November 2003 802.1 plenary meeting.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 34

See the disposition of comment 37.

Comment 35 Dinesh Mohan

NAME: Dinesh Mohan
COMMENT TYPE: TR
CLAUSE: 15
PAGE: 84
LINE:
COMMENT START:
Figure 15-1 does not represent the complete dual-relay model for two provider bridges represented in figure, which depict the customer-provider interface or UNI. This issue has been covered by Paul Bottorff and Steve Haddock in past meetings and during last meeting, we had generally agreed to include Paul's modified figures which would also allow us to align with MEF architectural constructs.
COMMENT END:
SUGGESTED CHANGES START:
Replace Figure 15-1 with Paul's figure titled "IEEE Peer relationship with Provider Bridges" that was presented during last IEEE meeting.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 35

See the disposition of comment 37.

Comment 36 Paul Bottorff

NAME: Paul A. Bottorff
COMMENT TYPE: T
CLAUSE: 15
PAGE: 84 LINE:
COMMENT START:

Figure 15-1 is confusing since it does not depict the customer handoff point.
COMMENT END:
SUGGESTED CHANGES START:
Replace Figure 15-1 with the attached slide.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 36

See the disposition of comment 37.

Comment 37 Muneyoshi Suzuki

NAME: Muneyoshi Suzuki
COMMENT TYPE: TR
CLAUSE: 15
PAGE: 84 LINE: Figure 15-1
COMMENT START:
Figure 15-1 seems to me inaccurate. As I addressed in the first comment, LLC in the C-VLAN aware Bridge is not located on the MIF. And there are no higher layer entities in S-VLAN aware Bridge.
COMMENT END:
SUGGESTED CHANGES START:
Fig. 3 in the attached ppt file illustrates proposed basic position of the Provider Bridge supported Service VLAN aware Bridging function within the MAC sublayer. Note that S-VLAN aware Bridges that locate provider edge don't need to support higher layer entities, because from viewpoint of Provider Bridged Network, provider edge behaves as an end station. Also note that this figure shows port-based service, but if service multiplexing is required, as proposed in the meeting, conjunction of C-VLAN aware and S-VLAN aware Bridges is needed as shown in Fig. 4.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 37

Discuss.

The comments all suggest making the figure show much more than was intended. The name given to the Figure clearly doesn't work very well, which is the editor's fault. The editor still wants to have a figure here, which is conceptually very similar to Figure 6-1. The principle purpose of figure 6-1 is to show that the relaying function of a bridge occurs below the MAC Service boundary, i.e. is transparent to MAC Service users in end stations. The purpose of Figure 15-1 is simply to show that, despite the bridging being carried out, the service provided by a provider bridges (S-bridges) in a provider bridge network is transparent to the operation of the customer VLAN-aware tagging functions in C-bridges. That's all. The editor believes that there needs to be a figure that simply conveys that, and welcomes new names for the existing figure. Replacing the S-Bridge to C-Bridge connection by an imaginary MAC just complicates things - the picture is true independent of whether there is a real or imaginary MAC between C-bridge and S-bridge. Similarly showing service interface points is also a distraction, the only service points being really relevant to the intent of the figure being the ISS SAPs below the C-TAG functions.

The editor welcomes new names for the existing figure, or simple additional text that points out the very specific and limited point it is trying to make. Other figures at other points may be required to meet the objectives of the comments.

3.21 Provider Bridge Description

Comment 38 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: T

CLAUSE: 15.4
PAGE: 86
LINE:
COMMENT START:
The figure should be modified to make the C-VLAN aware component peering to the S-VLAN aware component. It should be noted that the only time the C-VLAN component exist is when the C-VLAN need to be mapped/translated to S-VLAN. Therefore, C-VLAN doesn't get carried over S-VLAN as shown in the figure.
COMMENT END:
SUGGESTED CHANGES START:
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 38

Reject. The C-VID has to be carried over the S-VLAN any time more than one C-VLAN is to be supported by the same service provider service instance. The C-VID can be omitted (null tagged on the link to the Customer Network Port) or not omitted depending on how the C-Bridge Port is configured. It is not even necessarily the case that both ends of the logical link be configured the same, desirable for sanity, but not necessary for correctness.

Comment 39 **Ali Sajassi**

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 15.4
PAGE: 86
LINE:
COMMENT START:
The clause 5.9 defined three types of ports:
a) Provider Edge Port
b) Customer Network Port and
c) Provider Network Port
This clause should make use of these port types in its description
COMMENT END:
SUGGESTED CHANGES START:
SUGGESTED CHANGES END

Editor's Proposed Disposition of Comment 39

Accept. Needs final resolution of comment 5.

3.22 OAM loop detection

Comment 40 Muneyoshi Suzuki

NAME: Muneyoshi Suzuki
COMMENT TYPE: TR
CLAUSE: 16.8
PAGE: 95 LINE: 6-9 COMMENT START:
The following comment is a repeat of comment 46 in the disposition of ballot comments on P802.1ad/D1. This is because, while the disposition of the comment was "discuss", it has not addressed P802.1ad ballot resolution session in the meetings. As I presented in the last plenary meeting, Providers definitely need OAM that detects customer loops through Provider Bridged Network. Therefore, it seems to me, the assertion of "a network provider cannot rely on any customer network relaying such frames" in this clause is too strong. A customer network can selectively discard frames received from the provider. Therefore, there is no perfect solution that can detect connectivity loops through attached network. However, this does not mean providers can't rely capability of customer networks.

COMMENT END:
SUGGESTED CHANGES START:
Replace this sentence as following: A network provider may rely on customer network relaying such frames, but should develop an OAM and mechanisms to deal with potential data loops that can arise if the attached customer systems do not correctly operate its own instance or instances of spanning tree protocol.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 40

See disposition of comment 41.

Comment 41 Parsons, Bottorff, Mohan (OAM loop detection)

NAME: Glenn Parsons
COMMENT TYPE: T (Technical)
CLAUSE: 16.8
PAGE: 95 LINE:
COMMENT START:
Providers need the ability to detect customer loops given the fact that customer networks may or may not implement MSTP. They may rely on the customer (if they are 'trusted') or may develop their own mechanism. Both should be allowed.
COMMENT END:
SUGGESTED CHANGES START:
Reword 'cannot rely' to 'may rely'.
SUGGESTED CHANGES END:

NAME: Paul A. Bottorff
COMMENT TYPE: T
CLAUSE: 16.8
PAGE: 95 LINE:
COMMENT START:
An additional method for customer loop detection should be provided based on sourcing a provider OAM frame into the customer network.
COMMENT END: SUGGESTED
CHANGES START: Add a sentence to 16.8 indicating that provider loop detection OAM frames may also be used to detect customer loops.
SUGGESTED CHANGES END:

NAME: Dinesh Mohan
COMMENT TYPE: TR
CLAUSE: 16.8
PAGE: 95
LINE:
COMMENT START:
During last meeting, a proposal was made by NTT to detect customer loops using provider's loop detection OAM frames. It should be covered in this clause by making an additional note to this effect.
COMMENT END:
SUGGESTED CHANGES START:
Add an additional note (e.g. Note 5), to indicate that provider may rely on customer network to relay provider's OAM frames if detection of loops within the customer's networks is to be supported by provider.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 41

As circulated by email, and discussed in the January 2004 meeting. The editor would welcome further text if the existing NOTE is unsatisfactory, but will otherwise assume the comments are covered.

Clause 16.8 NOTE 3 "They can include ... customer agreement to allow the use of provider loop detection protocols." was added to D2.0 particularly to cover OAM loop detection. Since the scope of this amendment does not include normative requirements for OAM implementation, there isn't really a way to make a much stronger statement of permission.

The terminology of the standard is carefully chosen:

1) "X can rely on proposition P" means that proposition P is always TRUE for X, and "X cannot rely on proposition P" that P is not always TRUE for X. There is no relative "strength" in the statement. Either P is (always TRUE) or P is (NOT(always TRUE)).

2) The term "may" is strictly reserved in the standard for items which are options of the standard. On occasion we have not done a very good job of this in all our documents, but the guideline for production of the PICS has for many years been to collect up all instances of "shall" and "may" in the document. These terms are used as originally specified by ISO/IEC as they have approved translations between the major languages used by ISO. That means that "may" cannot be used as a general observation about behavior of something for which the standard does not make normative provisions. See the editors place holder note under 5.1 terminology. For there to be a "may" in this part of the discussion it has to apply to something which is in scope of implementation of the standard, and which is a choice to be made by the implementation.

If further improvement to the text is required, it is desirable that there be a proposal for new wording that:

a) allows the service provider to use OAM (the OAM to be specified elsewhere) e.g. "a service provider can use OAM to assist in the detection of loops created by customer networks" perhaps

(b) makes the difficulty clear "the customer can filter such frames"

(c) does not use "may", as 5.1 points out "may" and "may not" mean exactly the same thing (if it is optional to implement something, then it is equally optional not to implement it).

(d) incorporates the idea, as we discussed at the November meeting, that since the provider has to limit the amount of traffic from the customer, whether the traffic is caused by a loop or not, the detection of loops is of benefit to the customer (it is the customer's own bandwidth that is being wasted) therefore the customer ought to be as interested in helping detect loops, as the provider is in detecting them (provider is always interested in the customer perceiving the service to be of high quality).

3.23 Bandwidth limiting

Comment 42 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: T
CLAUSE: 16.8
PAGE: 95 LINE:
COMMENT START:
There is a lot of discussion in this clause of how bandwidth limitations can act as the solution to customer looping and provider service preservation. However, we have no specifications for how to do this, thus creating the 'slippery slope'. It would benefit this specification to consider including a simple specification for ingress bandwidth limiting.
COMMENT END:
SUGGESTED CHANGES START:
Consider adding ingress bandwidth limiting to this specification or consider the creation of a PAR of related documents that addresses this limited scope.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 42

Discuss. The editor would like a proposal.

3.24 Port and protocol VLANs

Comment 43 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: TR
CLAUSE: 16.3
PAGE: 91 LINE:
COMMENT START:
The first note isn't entirely true if port and protocol VLANs are supported on provider ports. In fact to job of supporting this classification type could be quite difficult and I believe further restrictions or clarifications are required.
COMMENT END:
SUGGESTED CHANGES START:
Consider dropping support for port and protocol VLANs on provider ports entirely. Alternatively, specify that they are only supported on customer network ports and on customer frames that are untagged originally. Parsing for protocol types on C-tagged frames would be complicated and resolution of conflicts would need to be explained. At a minimum, this note must clarify that for port and protocol VLANs (if supported) the frame is examined in this case.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 43

Accept.

Comment 44 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: TR
CLAUSE: 16.3
PAGE: 92 LINE: COMMENT START:
The last sentence of Note 9 is inconsistent with previous discussions supporting port and protocol VLANs. This sentence makes it sound as though port and protocol VLANs are not supported on provider ports.
COMMENT END:
SUGGESTED CHANGES START:
Make consistent with direction of port and protocol VLAN support
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 44

Accept. See also disposition of comment 43.

3.25 IETF reference

Comment 45 Glenn Parsons (IETF reference)

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 6.11
PAGE: 28 LINE:
COMMENT START:

The addition of this IETF dependency does not appear to add much value. That is both the 802.1ad work and IETF L2VPN work should be able to stand on their own without forward references between them. Any and all details of how provider bridging is mapped into foo should be defined by the group defining foo. As a result, nothing is needed in this standard.
COMMENT END:
SUGGESTED CHANGES START:
Delete 6.11
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 45

The comment assumes that the defining group for the protocol will define a satisfactory mapping from the ISS to the protocol, instead of just the protocol, which is counter to past experience. However since the editor now believes that there is sufficient understanding of where to put mappings if mappings are required (to ITU work if not to IETF, possibly), this clause can go until really needed.

Comment 46 Dinesh Mohan

NAME: Dinesh Mohan
COMMENT TYPE: TR
CLAUSE: 6.11
PAGE: 28
LINE:
COMMENT START:
As per comment on the previous version of .1ad draft, I find it difficult to understand the need to specify ISS by IETF foo in this specification. One could argue a similar clause for ITU-T foo. I believe that the focus should remain on current specification and leave the compliance work to respective standards bodies. Therefore specific reference to IETF should be removed.
COMMENT END:
SUGGESTED CHANGES START:
Delete clause 6.11
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 46

The comment assumes that the defining group for the protocol will define a satisfactory mapping from the ISS to the protocol, instead of just the protocol, which is counter to past experience. However since the editor now believes that there is sufficient understanding of where to put mappings if mappings are required (to ITU work if not to IETF, possibly), this clause can go until really needed.

Comment 47 Dirceu Cavendish

NAME: Dirceu Cavendish
COMMENT TYPE: TR
CLAUSE: 6.11
PAGE: 28 LINE:
COMMENT START:
Why is this clause needed? At most, it should be part of Clause 6.10
COMMENT END:
SUGGESTED CHANGES START:
Drop Clause 6.11
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 47

The comment assumes that the defining group for the protocol will define a satisfactory mapping from the ISS to the protocol, instead of just the protocol, which is counter to past experience. However since the

editor now believes that there is sufficient understanding of where to put mappings if mappings are required (to ITU work if not to IETF, possibly), this clause can go until really needed.

NAME: Paul A. Bottorff
COMMENT TYPE: T
CLAUSE: 6.11
PAGE: 28 LINE:
COMMENT START:
There is no need to reference IETF work to complete the 802.1ad standard. Adding this clause adds an un-needed inter-committee dependency.
COMMENT END: SUGGESTED CHANGES START:
Delete clause 6.11
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 47

The comment assumes that the defining group for the protocol will define a satisfactory mapping from the ISS to the protocol, instead of just the protocol, which is counter to past experience. However since the editor now believes that there is sufficient understanding of where to put mappings if mappings are required (to ITU work if not to IETF, possibly), this clause can go until really needed.

Comment 48 Dirceu Cavendish

NAME: Dirceu Cavendish
COMMENT TYPE: TR
CLAUSE: 15
PAGE: 84
LINE: NOTE
COMMENT START:
We shouldn't have a vague reference to MEF or ITU in the text.
COMMENT END:
SUGGESTED CHANGES START:
Remove last sentence of NOTE.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 48

Could someone (Dirceu?) provide specific references for the Bibliography, since we have agreed we should reference and not repeat or summarize the MEF/ITU work.

3.26 Maximum frame size 802.3 liaison

Comment 49 Paul Bottorff

NAME: Paul A. Bottorff
COMMENT TYPE: TR
CLAUSE: 1.0 PAGE: 1 LINE: 1
COMMENT START:
A general problem with Provider Bridging is that the 802.3 frame size needs to increased to support the Provider Bridge frame format. The 802.1ad committee needs to provide the requirements for 802.1ad frame extension to 802.3.
COMMENT END:
SUGGESTED CHANGES START:
Add footnotes indicating the need for co-ordination with 802.3
SUGGESTED CHANGES END

Editor's Proposed Disposition of Comment 49

Discussion is planned.

Comment 50 Glenn Parsons

NAME: Glenn Parsons
COMMENT TYPE: TR (Technical, Required)
CLAUSE: 9.7
PAGE: 76
LINE:
COMMENT START:
Given the addition of at least the S-TAG, all MACs that support this standard will require a larger frame size.
COMMENT END:
SUGGESTED CHANGES START:
A new PAR needs to be started in parallel in the appropriate 802 groups to ensure that the underlying MAC supports the larger frame size required by this standard to support provider bridging.
802.1 should approach at least 802.3 to request a PAR to support the frame growth required by 802.1ad
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 50

See disposition of comment 49.

Comment 51 Dinesh Mohan

NAME: Dinesh Mohan
COMMENT TYPE: TR
CLAUSE: General
PAGE:
LINE:
COMMENT START:
The resulting increased frame size due to S-VLAN in Provider Bridges requires co-ordination with 802.3. It is recommended that this coordination be initiated early enough to expedite the specification, when it gets completed.
COMMENT END:
SUGGESTED CHANGES START:
A comment should be maintained in this document indicating requirement for co-ordination and progress of such co-ordination with 802.3
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 51

See disposition of comment 49.

3.27 Typos etc.

Comment 52 Dirceu Cavendish

NAME: Dirceu Cavendish
COMMENT TYPE: TR
CLAUSE: All
PAGE: All LINE: All
COMMENT START:

Clauses 5, 8, 15 and 16 don't seem to be in synch, perhaps due to text from various generations of the amendment. For instance, clause 5 describes Customer-VLAN and Service-VLAN aware bridge components, which are later used in Clause 15 to define what a provider bridge is. The first figure that defines a provider bridge is Fig. 15-4. Clause 8(.6.1) talks about a VID translation table, however it is unclear which component (s) at Clause 5 realizes this translation. Clause 5 needs clarification text about VID translation table.
COMMENT END:
SUGGESTED CHANGES START:
Clause 5 needs clarification text about VID translation table.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 52

See editor's notes accompanying the ballot, and included in the front matter.

Comment 53 Dinesh Mohan

NAME: Dinesh Mohan
COMMENT TYPE: ER
CLAUSE: 1.1
PAGE: 3 LINE:
COMMENT START:
In (r), the term "user" should be changed to "customer" to maintain consistent use of customer terminology. It is also recommended that the same change be applied to remaining document, though other instances are not being pointed out.
COMMENT END:
SUGGESTED CHANGES START:
Change "user" to "customer".
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 53

Accept.

Comment 54 Dinesh Mohan

NAME: Dinesh Mohan
COMMENT TYPE: ER
CLAUSE: 5.7.1
PAGE: 12 LINE:
COMMENT START:
Numbering of bullets need to be corrected. Also move second (b) to "shall not" category.
COMMENT END:
SUGGESTED CHANGES START:
Numbering of bullets need to be corrected. Also move second (b) to "shall not" category.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 54

Accept.

Comment 55 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: E
CLAUSE: 5.7 PAGE: 12
LINE:
COMMENT START: <Details of comment> COMMENT END:
SUGGESTED CHANGES START:
In the first sentence, change "(C-TAG)" to "(S-TAG)".
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 55

Accept.

Comment 56 Steve Haddock

NAME: Stephen Haddock
COMMENT TYPE: E
CLAUSE: 6.3
PAGE: 25 LINE:
COMMENT START: <Details of comment> COMMENT END:
SUGGESTED CHANGES START:
In the last sentence on page 25, change "Table 6-3" to "Table 6-4".
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 56

Accept.

3.28 Other

Comment 57 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 15.5
PAGE: 87
LINE:
COMMENT START:
It is not clear to me what the clause intend to say by saying that "... Provider Bridges operated by a customer ..." since they are never operated by customers.
COMMENT END:
SUGGESTED CHANGES START:
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 57

Reject. There is no embargo on the customer acquiring and operating any piece of equipment. The term Provider Bridge is a convenience since bridges with this functionality are usually used by providers, b=ut that does not prevent their use by the customer or creation of a service that relies upon their use by some customer (a carrier's carrier supporting one of its customers for example).

Comment 58 Paul Congdon

NAME: Paul Congdon
COMMENT TYPE: E
CLAUSE: 16.3
PAGE: 92 LINE:
COMMENT START:
Note 8 would make more sense if it followed the paragraph on the previous page that indicates egress filtering is enabled always for ingress ports.
COMMENT END:
SUGGESTED CHANGES START:
Move note 8 here to keep discussion of ingress filters relevant.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 58

Accept.

Comment 59 Arjan de Heer

NAME: Arjan de Heer
COMMENT TYPE: T
CLAUSE: 15.5
PAGE: 87 LINE:
COMMENT START:
This describes the scenario of interconnection between two providers, it is a kind of 'NNI' instead of 'UNI'. I am not in favour of having customers adding the tag. The tag is fully owned and managed by the provider.
COMMENT END:
SUGGESTED CHANGES START:
Adapt text.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 59

Reject. The document shouldn't prevent providers creating such a service if they want to. It does not necessarily imply a peer relationship. See disposition of comment 57.

Comment 60 Arjan de Heer

NAME: Arjan de Heer
COMMENT TYPE: T
CLAUSE: 16.3
PAGE: 92 LINE:
COMMENT START:
Use of VID translation table for provider-provider scenarios, as in Note 6. I would rather have the scenario where a user used the S-tag as a note, and the provider interconnect as the main text. From the text it is not clear whether the VID translation table only translates in ingress direction, or in both ingress and egress direction. First sentence on page 92 talks about ingress Ports, but Note 6 talks about ingress and egress direction. 8.6.1 ingress is mentioning it, but 8.6.4 is not. Furthermore, apart from a VID translation table a Priority translation table is required, because two providers may use different encoding schemes. (Using a priority regeneration table, but that is prohibited for tagged frames)
COMMENT END:
SUGGESTED CHANGES START:
Adapt text, clarify VID translation table, add Priority translation
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 60

Accept in principle. Needs to be clear that both ingress and egress translation occur. The ideal is that one VLAN has one VID "on the wire" (excluding shared learning scenarios). See also the discussion on priority translation, which does assume that the priority coding is consistent for all S-VLANs on a port.

Comment 61 Arjan de Heer

NAME: Arjan de Heer
COMMENT TYPE: T
CLAUSE: 16.2
PAGE: 91 LINE:
COMMENT START:
Suppose we have a provider bridge, with a C-VLAN aware component with two Provider Edge ports. Both provider edge ports are a member of C-VLAN 1033. Is this configuration excluded by item b and c? If so this should be made explicit.
COMMENT END:
SUGGESTED CHANGES START:
Give clarification
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 61

Discuss. The editor has prepared two alternative responses to this comment:

1) There is no intent to exclude such a configuration. In fact it is highly desirable that it not be excluded, because it can be used to provide elements of reliability through redundancy. The necessary text change can be effected by replacing "within the provider networks" with "within the network composed of the S-VLAN aware bridge components of the provider's Provider Bridges and Provider Edge Bridges". This is a clear illustration of the fact that service selection and segregation policing has to be performed at the Customer Network Ports of the S-bridge components, even when they are not the ports that the customer connects to with an access LAN.

2) While two separate items of equipment can be configured as suggested, there appears to be no benefit to making the composition of a Provider Edge Bridge so flexible. The C-VLAN components of a Provider Edge Bridge should be specified to be one per Provider Edge Port.

Comment 62 **Ali Sajassi**

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 16.9
PAGE: 95 LINE:
COMMENT START:
The terminology defined in here is inconsistent with the ones defined earlier - e.g., Provider Network Port, Provider Interface Port, Customer Interface Port, etc.
COMMENT END:
SUGGESTED CHANGES START:
Remove PIP, CIP and use Provider Network Port and Provider Edge Port respectively
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 62

Accept in principle.

Comment 63 **Ali Sajassi**

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 16.9
PAGE: 95
LINE:
COMMENT START:
There is no reason as to why a CE cannot be directly connected to a PPE. Change the last paragraph to:
COMMENT END:
SUGGESTED CHANGES START:
The application of the ingress and egress rules, together with the use of the MSTP restrictedRole and GVRP registration controls (16.4), permit providers to allow direct attachment of customer operated equipment to access LANs connected to PPE PIPs. Locating provider bridging functions within CPE, as illustrated by but not limited to the examples CPE1 through CPE3 in Figure 16-1, can be used to partition and enhance provider network access functionality in support of:
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 63

Discuss. Many providers would disagree on maintenance and failure isolation grounds.

Comment 64 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 16.9
PAGE: 96
LINE:
COMMENT START:
Change the Figure 16-1 to show a CE is directly connected to a PPE.
COMMENT END:
SUGGESTED CHANGES START:
SUGGESTED CHANGES END

Editor's Proposed Disposition of Comment 64

See disposition of comment 63.

Comment 65 Ali Sajassi

NAME: Ali Sajassi
COMMENT TYPE: T
CLAUSE: 16.9
PAGE: 96
LINE:
COMMENT START:
I don't see why we need a CPE in order to satisfy (b), (c), (d), and (f). I think we really don't need to make a distinction between CPE and PPE in terms of functionality. They can both have equal functionality (e.g., both can include C-VLAN aware and S-VLAN aware components). So I would suggest to change this clause to reflect this.
COMMENT END:
SUGGESTED CHANGES START:
SUGGESTED CHANGES END

Editor's Proposed Disposition of Comment 65

See disposition of comment 63.

Comment 66 Mick Seaman

NAME: Mick Seaman
COMMENT TYPE: E
CLAUSE: All
PAGE: All LINE:
COMMENT START:
Given the extremely low number of responses to the ballot, and the few comments on detail as opposed to the higher principles that remain in dispute. I can only recommend that the document goes forward to working group status so that eventually non-respondents will lose voting status.
COMMENT END:
SUGGESTED CHANGES START:
Resolve the comments that disagree with the initial scope as stated at PAR approval time, i.e. no service interface modifications and no new major entities, by voting at the meeting on the intention to progress new PARS (or not) to address these extensions. Determine at the meeting the technical basis for requiring VID space extensions to support the required number of provider service instances, and vote on the optional inclusion (or not) of such a VID. This is a necessary precursor to communication with 802.3 on the subject of larger frame sizes. Following resolution of these major issues, progress the document to working group level.
SUGGESTED CHANGES END:

Editor's Proposed Disposition of Comment 66

Accept :-)