

IEEE P802.1Qau D2.1 Congestion Notification comments

CI 03 SC 3.2 P4 L9 # 15  
 Mick Seaman None entered

Comment Type TR Comment Status A

The text of this definition goes beyond that required for a definition, and starts to explain normative provisions of the standard. Such additional explanation is at best duplication, and is likely to conflict. The definition needs to be trimmed to the necessary and accurate. This definition is inaccurate because it states that frames are no longer part of a CCF when they are not in the end station: "Outgoing frames in an end station ..". The definition should not go on to expand on the relationship between Reaction Points and Flow queues as stating "the same Flow queue and Reaction Point" suggests that a single Reaction Point could handle multiple "Flow queues" which is not the case. Similarly the text currently implies that different priority values could be assigned to the same Flow queue. "Flow" is not a reserved word and hence should not be capitalised.

*SuggestedRemedy*

Replace the existing definition with the following: "A sequence of frames with the same priority, that the transmitting end station treats as belonging to a single flow, using a single Reaction Point to controlling the transmission rate of all those frames."

It should be apparent from this definition (and is a little obscure in the original) that the efficacy of CN will depend somewhat on the degree to which the frames in a CCF traverse the same CPs.

If absolutely necessary to satisfy past comments the additional sentence "Multiple CCFs can be assigned to the same Reaction Point." could be added but would be better dealt with in the definition of Reaction Point.

Response Response Status C

ACCEPT. (Sentence about multiple CCFs per RP is included in suggested remedy of Comment #20.)

CI 03 SC 3.3 P4 L13 # 16  
 Mick Seaman None entered

Comment Type TR Comment Status A

This definition would be a little clearer if it were edited to be briefer. By describing the conditions under which a given priority is or is not a CNPV it goes a little too far, and is apparently technically incorrect. I don't believe that two ports of a bridging system could be configured to use a given priority as a CNPV, while two other ports are not. It seems inevitable that a system should treat a given priority as a CNPV (or not) on all ports, even if the domain defense mode is different on different ports. It is probably not worth extending this to "domain" within this definition because that would overlap or recurse with the CNP definition. As minor point, I believe that "port" should not be capitalised (it should only be capitalised where the term is being used as part of a specific defined term, not in general use.)

*SuggestedRemedy*

Replace the existing definition with the following: "

A value of the priority parameter that a congestion aware system uses to support congestion notification."

Response Response Status C

ACCEPT.

CI 03 SC 3.5 P4 L23 # 17  
 Mick Seaman None entered

Comment Type TR Comment Status A

This definition repeats itself, but the duplicated information is not entirely consistent and not accurate in either case: "A message indicating the degree of congestion of the queue" is almost a duplicate of "carries information that indicates to what degree the .. is to reduce the rate" but neither is really true. The first statement leaves out the derivative, the second taken literally would mean that the CP would have to have information about the RP's rate. The potential location of CP's is dealt with in the CP definition and does not need to be repeated here. Similarly the location of RP's is dealt with in the RP definition.

*SuggestedRemedy*

Replace the existing definition with the following: "

A message transmitted by a Congestion Point to a Reaction Point, in response to a frame received from that Reaction Point, conveying congestion information used by the Reaction Point to reduce its transmission rate."

Response Response Status C

ACCEPT.

IEEE P802.1Qau D2.1 Congestion Notification comments

CI 03 SC 3.6 P4 L29 # 18  
 Mick Seaman None entered

Comment Type TR Comment Status A

"that can be inserted into every frame" does not appear to make strict sense. Does it mean in every frame or in no frames at all? If "every" then more conditions and explanation would be required (of CCFs, CNPVs etc.). The sentence has further difficulties since it is not clear whether the tag, the frame, or values from the tag, or the flow identifier can or is to be returned in a Congestion Notification Message. The sentence, with its attempt to repeat (or modify) provisions of the standard, goes beyond that is required for a definition and repeats information from other definitions that are only a few lines away. This repetition should be removed.

SuggestedRemedy

Replace the definition of a Congestion Notification Tag with the following:  
 A tag that conveys a Flow Identifier, that a Reaction Point can add to transmitted CCF frames.

Response Response Status C

ACCEPT IN PRINCIPLE. Suggested wording is an improvement. However, the CN-tag is always returned at the head of a CNM. Will use:

A tag that conveys a Flow Identifier, that a Reaction Point can add to transmitted CCF frames, and that a Congestion Point includes in a Congestion Notification Message.

CI 03 SC 3.8 P4 L36 # 19  
 Mick Seaman None entered

Comment Type TR Comment Status A

The fact that the Flow Identifier can be used in a CN-TAG is said elsewhere, but the important properties of the Flow Identifier are missing.

SuggestedRemedy

Replace the definition of a Flow Identifier with the following:  
 An identifier assigned by a congestion aware end station, unique within the scope of one or more of the source MAC addresses used by that system to transmit CCF frames, that can be used to associate each received Congestion Notification Message with the Reaction Point that rate controls the CCF that caused its transmission."

Response Response Status C

ACCEPT.

CI 03 SC 3.9 P4 L42 # 20  
 Mick Seaman None entered

Comment Type TR Comment Status A

The construction of this definition leaves the scope of the "optionally" unclear. There is no real need to describe whether CN-TAGs are added or not within the scope of this definition. The real defining attribute is the rate control applied. Discussing the "reception or absence" of Congestion Notification Messages is a not very helpful trivialisation of the congestion control algorithm.

SuggestedRemedy

Replace the definition of Reaction Point with the following:  
 An end station port function (<ref>) that controls the transmission rate of frames for one or more CCFs, receiving and using Congestion Notification Messages as part of determining that rate."

Response Response Status C

ACCEPT.

CI 05 SC 5.4.1.4 P6 L10 # 21  
 Mick Seaman None entered

Comment Type ER Comment Status A

This clause is tucked one level too low in the clause numbering scheme. This is bad policy as it makes clause numbering cumbersome. Put the congestion notification requirements at the 5.4.n level (see 802.1aq Conformance clause for an example).

SuggestedRemedy

As per comment.

Response Response Status C

ACCEPT.

CI 05 SC 5.4.1.4 P6 L22 # 22  
 Mick Seaman None entered

Comment Type TR Comment Status A

The phrase "required specifications" is not well defined. If what is meant is specification required by this standard (802.1au) then these need to be identified, by reference. If what is meant is conformance to the mandatory (as opposed to any optional) provision of 802.1AB then that should be said.

SuggestedRemedy

As per comment.

Response Response Status C

ACCEPT IN PRINCIPLE. Change "specifications" to "capabilities" and add "-2005 Clause 5.2" after 802.1AB.

IEEE P802.1Qau D2.1 Congestion Notification comments

Cl 31 SC 31.2 P69 L49 # 23  
 Mick Seaman None entered

Comment Type TR Comment Status A

The sentence "The service to higher layers is not an instance of the ISS .." is pure bilge, and clearly results from a misunderstanding of what the ISS is and is not. Refer to 802.1Q Clause 6.1 and in particular to the sentence is 6.1.2 "The parameters of service primitives do not include information that is used only locally .." The point is that the ISS (or any reference model service is not an API) -- it contains only those elements necessary to express communication between peer systems. HOWEVER it is still necessary for there to be peers, even if one of the peers is within an end station. We cannot afford to degenerate into the sort of miscellaneous hackery that will result from throwing the notion of communicating peers away. Not only should we describe the upper boundary of the shim discussed as providing one or more instances of the ISS, it is also perfectly possible to add the local flow control functions necessary to add those (local) API characteristics to the ISS - the LMI (Layer Management Interface) provides all the flexibility and abstraction required -- clearly since the identification of CCFs and other aspects of end station behavior are not constrained by this standard it is not actually required that the LMI parameters be defined -- it is enough that they could be.

*SuggestedRemedy*

Reword the sentence to state that one (or more) instances of the ISS are being provided. If necessary explain that the ISS itself comprises only those parameters that express peer to peer communication, and exclude purely local issues (like internal flow control) though the latter can be modeled within the same framework as the ISS through the specification of LMI parameters (to support the exchange of local management information between layers).

Response Response Status C

ACCEPT.

Cl 00 SC 0 P0 L0 # 24  
 Mick Seaman None entered

Comment Type TR Comment Status R

Despite the changes in this draft that explain the use of CN with PBBNs (a very significant use case) I remain unconvinced that the supposed additional throughput to be realised by having multiple RPs per port will actually be of meaningful benefit. Moreover the suggestion that the Flow Id is actually a covert extension to MAC Addressing for the purpose of identifying individual links from a link aggregation in order to get CNMs to a physical interface shows that we are going to get ourselves into miscellaneous hackery in that most dangerous of areas, addressing.

*SuggestedRemedy*

Allow only one RP per port (i.e. per end station MAC address in cases where the meaning of port may be disputed), remove all the provision for CN-tag, and simplify the document throughout accordingly.

Response Response Status U

REJECT. There is a consensus that the CN-tag is necessary for some use cases, but not for others, and that it should be optional. This comment will be held over in Annex Z.