

CI 00 SC P L # 31
Thompson, Geoff GraCaSI S.A.

Comment Type TR Comment Status R

This project has failed to live up to the level of participation that was advertised in the PAR: "5.1 Approximate number of people expected to be actively involved in the development of this project: 30"

and it would appear that its market projections as put forth in the BMP criterium were overly optimistic on a grand scale. This is show by the poor participation. It would appear that most of the current interest comes from a particular industrial sector which failed to follow 802.3 recommendations about 20 years ago and did not install 4-pair cabling. Participation by other sectors has been very poor. IF there ever will be a true market need for this standard, it should be developed with broad participation from the bodies who need it when their own need is sufficiently close that the affected parties will send participants who are in the midst of development. It is a bad idea to develop a standard before the market is read for it.

SuggestedRemedy

Withdraw the project at this time or hibernate it until more people who are willing to participate in its development show up in 802.3. Requalify it for Broad Market Potential at that time and modify the PAR if needed and it is still active.

Response Response Status U

REJECT. The market projections in the Broad Market Potential based on the automotive and industrial environments continue to be accurate. In fact, there is interest in additional markets such as carrier backhaul and professional audio video.

We have active participation in joint meetings from IEEE 802.1 TSN (a group of more than 30) which has a companion project (IEEE P802.1Qbu Frame Preemption) dependent on this project. Also, about 30 people have participated by commenting on ballots.

The interest in operating on fewer pairs and at lower speeds in the automotive and industrial market is driven by the need to reduce weight and power consumption.

CI 00 SC 0 P L # 13
Trowbridge, Steve Alcatel-Lucent

Comment Type TR Comment Status R

The terminology in the amendment does not match the agreed objectives for the project. The Call for Interest held in the March 2012 plenary for Frame Preemption was withdrawn after too much controversy over the characterization of the problem and solution. After a subsequent CFI, the first attempt to approve a PAR and objectives at the July 2013 plenary in Geneva failed due to inconsistency of the terminology with 802.3 (distinguished minimum latency traffic and "M-frames", "M-frames in the wild" were rejected. After rework in the York interim, a characterization as "interspersing express traffic" was developed, leading to the currently accepted objectives accepted in November 2013. The only place the accepted terminology appears in the draft is in the title and the name of the task force. The entire draft uses the terminology of the withdrawn CFI from March 2012

SuggestedRemedy

Update the terminology globally in the draft per the agreed objectives. In particular:

1.4.3 - change "preemptable Media Access Control" to "non-express Media Access Control" with an appropriate acronym

1.4.4 - change "preemptable traffic" to "non-express traffic"

Add IET to the acronyms defined in clause 1.

Occurrences of "preemptable" in clause 30 change to "non-express", objects such as "PreemptSupported", "PreemptEnabled", "PreemptActive" change to "IETSupported", "IETEnabled", "IETActive", etc.

Change "preemption capability" to "IET capability" globally in clause 79.

pMAC and PMAC not consistent in clause 79, but should change globally to neMAC (or whatever acronym is chosen for the non-express MAC).

Clause 99: preemptable MAC should be non-express MAC globally.

"MAC client supporting preemption" becomes "MAC client supporting IET" globally.

pMAC becomes neMAC (or chosen acronym) globally

"preemption is active" becomes "IET is active" globally

"enable preemption" becomes "enable IET" globally

"link partner supports preemption" becomes "link partner supports IET"

Response Response Status U

REJECT.

The main complaint about the intial CFI was that it presumed a solution and that should be decided after the project is created.

After the project was created, preemption was chosen as part of the solution for interspersing express traffic. The suggested name changes would not aid the reader in understanding the material. There is no reason to obfuscate the selected mechanism.

The project meets the agreed objectives.

CI 90 SC 90.4.3.1.1 P 32 L 21 # 29
Thompson, Geoff GraCaSI S.A.

Comment Type TR Comment Status R

This addition of another variable seems unnecessarily complex. The bridge (or end station) is supposed to have port configuration information that knows this is a pMAC and therefore unsuitable for use in timed applications. Second, the indication should only take place upon the passage of a legacy SFD. The new SFD codings will not exert it.

SuggestedRemedy

Removed the new text.

Response Response Status U

REJECT. We asked IEEE 802.1 TSN at our joint meeting in July whether they needed this on the preemptable path or whether they could work with it only on the express path. The experts there affirmed that they need the time stamp on both paths.

CI 90 SC 90.4.3.2.1 P 32 L 43 # 30
Thompson, Geoff GraCaSI S.A.

Comment Type TR Comment Status R

New text is unnecessary

SuggestedRemedy

Remove new text.

Response Response Status U

REJECT. See #29

CI 99 SC 99.1 P 35 L 22 # 33
Dawe, Piers Mellanox

Comment Type TR Comment Status A

"the MAC Merge sublayer may prevent the pMAC from starting transmission of preemptable traffic." So this proposed thing is clearly a new MAC, because it controls access to the medium. A new MAC client with roughly twice as many queues, management registers, everything, is needed to use it. This isn't "Conformance with the IEEE Std 802.3 MAC", "conformance with the MAC client interface" or "conform to the full-duplex operating mode of the IEEE 802.3 MAC" as alleged in the 5C "Compatibility" response. It forces anyone with a MAC design to redesign it.

SuggestedRemedy

Revise the 5C responses to reflect that this is a new or modified MAC, get a vote from 802.3 as to whether they want that;
or revise the draft so that it conforms to the 5C "Compatibility" response;
or terminate the project, like P802.3ar Congestion Management.

Response Response Status U

ACCEPT IN PRINCIPLE. Replace with "the MAC Merge sublayer may prevent the start of transmission of frames from the pMAC"

It isn't changing the MAC. It is holding off acceptance of the primitive from the MAC. There is no change to the MAC. We are consistent with the Compatibility response since we do not make any changes to the MAC. Other projects such as PAUSE, PFC and point-to-multipoint changed the control of access to the medium without changing the MAC.

IEEE 802.1Qbu is defining protocols for MAC Clients that expect this behavior. It doesn't require twice as many queues. IEEE 802.1Q already defines use of up to 8 traffic classes (e.g. queues) and such implementations are common.

This is an optional capability and doesn't force anyone to support it. Devices supporting the optional capability are fully interoperable with devices that don't support it.

CI 99 SC 99.1 P 35 L 46 # 26
Thompson, Geoff GraCaSI S.A.

Comment Type ER Comment Status A

The definition of "conjunction" [noun: the action or an instance of two or more events or things occurring at the same point in time or space.] doesn't really work here. Please redo the text.

SuggestedRemedy

I suggest the following: "A MAC Control Sublayer associated with an eMAC or a pMAC shall not generate PAUSE when the associated MAC Merge sublayer is active."

Response Response Status U

ACCEPT IN PRINCIPLE. There is no concept of the MAC Merge sublayer being active. It is instantiated or not.
"A MAC Control Sublayer that is the client of an eMAC or a pMAC shall not generate PAUSE."

CI 99 SC 99.2.2 P 39 L 5 # 27
Thompson, Geoff GraCaSI S.A.

Comment Type TR Comment Status R

I see no need for this primitive. If the merge function is enabled and a frame is presented to the eMAC for transmission then it should be transmitted ASAP and any necessary preemption should take place without any further control needed.
Any hold-off function needed on the pMAC side can take place at the transmit buffer in the bridge.

SuggestedRemedy

Remove sub-clause 99.22

Response Response Status U

REJECT. This primitive is required in the project objectives.
"Provide a primitive at the MAC client service interface to inhibit the transmission of non-express frames."
This primitive allows the MAC client to preempt before scheduled traffic is due to arrive so that the scheduled traffic can be sent immediately. That cannot be done efficiently in buffer above the MAC because that would require stopping transmission a before the frame starts on the pMAC wasting up to a max frame time on the media. See the July 2013 Geneva Tutorial on IET, slide 39 for an illustration on this.

CI 99 SC 99.3.3 P 40 L 37 # 28
Thompson, Geoff GraCaSI S.A.

Comment Type TR Comment Status R

I am pretty unhappy with the entire approach of having multiple new values of the start frame/packet delimiter. To my knowledge there has been no investigation of the error robustness of such a scheme, especially one with multiple values. At the time of the initial approval there was significant discourse and investigation of the error robustness of the SFD. One of the results of that discussion was to require additional error checking on a per packet basis by the addition of a length field.

SuggestedRemedy

Use a scheme that doesn't require a new frame delimiter or delimiters. Using the established delimiter will at least provide equivalent performance to current implementations.

Response Response Status U

REJECT. The existing delimiter has zero Hamming distance (a 1 bit change during preamble can cause a false SFD). The new delimiters all have at least a Hamming distance of 3 from preamble (and a Hamming distance of 4 from SFD). They are therefore all stronger than the original SFD and have the same Hamming distance from SFD that we have used in developing PHY encodings such as 64b/66b.

While IEEE 802.3 initially added a length field to strengthen the SFD, most frames today use an Ethertype and IEEE 802.3 was updated to allow that.

CI 99 SC 99.3.3 P 40 L 37 # 23
Thompson, Geoff GraCaSI S.A.

Comment Type TR Comment Status R

Changing delimiters means that all media side test equipment for this (small market) technology will have to have a hardware change from legacy equipment. If a scheme were used that kept the legacy delimiter, then legacy and current main market test equipment could be used in IET applications with only a software change

SuggestedRemedy

Use a scheme that doesn't require a new frame delimiter or delimiters. Using the established delimiter will at least provide hardware compatibility with broad market test equipment both in manufacturing and in the user field.

Response Response Status U

REJECT. All mechanisms that don't introduce new delimiters require additional overhead for added headers. This mechanism was chosen because it adds no additional overhead to unpreempted frames and minimizes the overhead for preempted frames to the extent possible while meeting other objectives.

In addition, this does not require a change to all test equipment. Some test equipment captures the full packet including preamble, has programmable SFD capture or other mechanisms that don't require hardware change.