Unsatisfied comments from initial ballot as of 16 July 2015

IEEE P802.3br (D2.0) Interspersing Express Traffic Initial Working Group ballot comments

Cl **00** SC P L # 380

Peter Stassar Huawei Technologies

Comment Type **ER** *Comment Status* **A** *Preepmt vs IET*

The draft is totally inconsistent between its title, referring to Interspersing Traffic and the actual text, where only 'Preempt ...' is being used.

SuggestedRemedy

Fix inconsistency.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE. Preemption is the mechanism that allows for interspersing express traffic. Add to 99.1 to explain the relationship.

The MAC Merge sublayer supports interspersing express traffic with preemptable traffic. This is achieved by using a MAC Merge sublayer to attach an express Media Access Control (MAC) and a preemptable MAC to a single Physical
Signaling Sublayer (PLS) service. The MAC Merge sublayer supports this with two methods to stop transmission of preemptable traffic so that express traffic can be transmitted. It can preempt or not initiate transmission of preemptable traffic so that express traffic can be transmitted.

Cl **00** SC **0** P L # 331

Trowbridge, Steve Alcatel-Lucent

Comment Type **TR** *Comment Status* **A** *Preepmt vs IET*

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 The Call for Interest held in the March 2012 plenary for Frame Preemption was withdrawn after too much controversy over the characterization of the problem and solution.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 The entire draft uses the terminology of the withdrawn CFI from March 2012.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**

ACCEPT IN PRINCIPLE. The terminology was agreed to in the base line proposal by the task force.

Preemption is the capability that provides for interspersing express traffic.

See also #380 for some changes to better relate the two terms.

Cl **00** *SC* **0** *P* *L* # 58

Grow, Robert — RMG Consulting

Comment Type **TR** *Comment Status* **R**

I am unable to convince myself that the amendment doesn't make what is to me are unacceptable and unstated assumptions of compatible MAC and PHY characteristics. For example, if it assumes all PCS layers use codes that either encode less than an octet (e.g., Manchester bit encoding) or that have an integer number of octets in the PCS code. This is a new requirement. I did not find a requirement that mPackets had to be contiguous and could not cause interframe to be signaled on an xMII unless until both a pFrame and one or more eFrames are completely transmitted when a preemption occurs. Failure to do this could result in RX_DV being deasserted falsely indicating an end of frame on the xMII.

I believe this is a problem for PCS layers that do not encode an integer number of octets. For example, if a 10 Mb/s or 100BASE-X MAC produces a non-integer number of octets, the MII nd currently defined PHYs convey that across the link so that an alignment error can be detected.

I similarly worry that a PHY code that does not include an integer number of octets in a code word could result in a false indication of interframe spacing at the receive xMII.

SuggestedRemedy

Assure MAC Merge will properly convey an alignment error across a link and that contiguous mPackets are required so that interframe will not be improperly created at a receive xMII.

Response *Response Status* **U**

REJECT.
Receive processing receives the packet a bit at a time and does not assume that it is an integer number of octets in length.

There is no assumption that mPackets are contiguous. They must be separated by at least an interpacket gap.

Cl **00** *SC* **0** *P* *L* # 57

Grow, Robert — RMG Consulting

Comment Type **TR** *Comment Status* **A** *PAUSE*

Other than Figure 99-1, and a few other mentions of MAC control as part of express traffic delay requirements, the amendment doesn't address interaction with MAC Control pause. It seems that impacts on pause quanta and interruptability of MAC control frames should be addressed. Were these other optional protocols considered in development of this amendment?

SuggestedRemedy

Please address.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE. Interoperation with MAC Control PAUSE and PFC was considered.

Add to 99.1: "A MAC Control Sublayer shall not generate PAUSE when used in conjunction with MAC Merge."
PAUSE would only affect the MAC Control sublayer on which it was received unless work was done to redefine how it worked with two MAC Control sublayers above two MAC Merge sublayers. It would make more sense to use PFC.

With PFC, IEEE 802.1Qbu should discuss the interoperation of PFC and preemption. This has been discussed with the TSN task group during our joint meetings. They are handling it in their draft which currently says to send PFC requests to the eMAC Client interface

Cl **01** *SC* **1.4.1** *P* **16** *L* **17** # 381

Thompson, Geoff — GraCaSI S.A.

Comment Type **TR** *Comment Status* **A**

The current text of the definition appears to require the definition of a "new MAC". My impression of this project was that it was supposed to accomplish its goals within the reconciliation sub-layer and use two instances of a normal full-duplex MAC.

SuggestedRemedy

Change text to read: "1.4.1 express Media Access Control (eMAC): The instance of the Media Access Control sublayer associated with an Interspersing Express Traffic port which is the client of a MAC Merge sublayer service interface that handles express frames."

Response *Response Status* **U**

ACCEPT IN PRINCIPLE. IEEE 802.3 does not use the term port except in a very limited sense (i.e. where a fiber optic cable attaches) so this definition wouldn't work.

"The instance of a Media Access Control sublayer (IEEE Std 802.3 Annex 4A) which is the client of a MAC Merge sublayer that handles express traffic."

Do the same for pMAC and preemptable traffic.

Cl **01** *SC* **1.4.5** *P* **14** *L* **27** # 69

Hajduczenia, Marek — Bright House Network

Comment Type **ER** *Comment Status* **A**

"See IEEE Std 802.3br, Clause 99." - we reference clauses, and not specific amendments.

SuggestedRemedy

Change to "See IEEE Std 802.3, Clause 99."

Response *Response Status* **U**

ACCEPT.

Cl **1** *SC* **1.4.3** *P* **16** *L* **22** # 382

Thompson, Geoff — GraCaSI S.A.

Comment Type **TR** *Comment Status* **A**

The current text of the definition appears to require the definition of a "new MAC". My impression of this project was that it was supposed to accomplish its goals within the reconciliation sub-layer and use two instances of a normal full-duplex MAC.

SuggestedRemedy

Change text to read: "1.4.3 express Media Access Control (eMAC): The instance of the Media Access Control sublayer associated with an Interspersing Express Traffic port which is the client of a MAC Merge sublayer service interface that handles preemptable frames."

Response *Response Status* **U**

ACCEPT IN PRINCIPLE. See #381

Cl **30** *SC* **30.12.3.1.27** *P* **20** *L* **19** # 91

Hajduczenia, Marek — Bright House Network

Comment Type **TR** *Comment Status* **R**

Attribute aLldpXdot3RemAddFragSize has very cryptic definiton: "A 2-bit integer value used to indicate, in units of 64 octets, the minimum number of octets over 64 octets required in non-final fragments by the receiver on the given port associated with the remote system;"

SuggestedRemedy

Is the intention to define the minimum fragment size? It would make much more sense to simply define it as INTEGER and then record the fragment size, and not some fragment size delta - these are MIB objects and not hardware registers!
Similar comment on aMACMergeAddFragSize

Response *Response Status* **U**

REJECT. All fragments have a minimum size of 64 octets. The purpose of this object is to request a size larger than that minimum for non-final fragments. If it was specifed as the fragment size rather than additional fragment size, we would have to define what happens for 0 which wouldn't be a legal minimum fragment size. By making it additional fragment size, there are no illegal values and each value means something distinct.

Cl **79** *SC* **79.3** *P* **24** *L* **14** # 93

Hajduczenia, Marek — Bright House Network

Comment Type **TR** *Comment Status* **A**

TBD in Table 79–1 - time to decide what this is going to be

SuggestedRemedy

Change TBD with the appropriate value for this new "Additional Ethernet Capabilities" subtype. The same value should be then propagated into 79.3.6 as well and Figure 79–6. "6" seems to be the next free number as of 802.3bx

Response *Response Status* **U**

ACCEPT IN PRINCIPLE. See #280
6 is in use by EEE.

Cl **99** *SC* *P* **32** *L* # 384

Thompson, Geoff — GraCaSI S.A.

Comment Type **TR** *Comment Status* **R**

This clause seems to (a) not precisely specify which configuration of the existing MAC is used for the eMAC and the pMAC and also seems to be respecifying the upper MAC service interface.

SuggestedRemedy

Respecify things so that the accommodation (and the accompanying implied buffering) take place in the MAC MERGE and RECONCILIATION sub-layers.

Response *Response Status* **U**

REJECT. It specifies that the MACs are full duplex operating at 100 Mb/s or greater (first line in 99.1). It is using two copies of the upper MAC service interface, not respecifying it. This was indicated as a example of how this might be implemented even before the PAR was approved.

Cl **99** *SC* **99** *P* **45** *L* **38** # 25

Anslow, Pete Ciena

Comment Type **TR** *Comment Status* **A** *Discuss*

There are several Editor's notes in Clause 99 discussing issues with the clause.
All of these issues should have been resolved prior to WG ballot and will certainly have to be resolved prior to the draft being ready for Sponsor Ballot.

SuggestedRemedy

Resolve all of the issues and remove the editor's notes.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE. There are 2 editor's notes that relate to issues. One documents a small issue in 30.14.1.2 that the editor noticed during draft preparation. There are comments that resolve this issue so this note should be gone in the next draft.

The other requests review of delay constraints (though the statement that it is a first cut is old and should have been removed - there has been some review and update during the task force review). This note will be removed in the next draft.

The other editor's notes are not on technical issues.
One highlights changes to the Containment diagram for voters because that was requested since the text change marking isn't in figures. Remove in the next draft.

Another provides an explanation of the value used for HRT. Remove in the next draft.

Cl **99** *SC* **99.3.3** *P* **36** *L* **49** # 385

Thompson, Geoff GraCaSI S.A.

Comment Type **TR** *Comment Status* **R**

I am opposed to the extent to which the SMD breaks the architecture of the long-standing Ethernet frame format and architecture by loading data content into the start frame delimiter.

SuggestedRemedy

Have only one new value of start frame delimiter whose job is to signal that the frame is a pre-temptable frame and handle all of the data for managing broken frames within the data field. I would strongly prefer that all such management data appear behind an EtherType field so things are consistent with other varieties of VLAN frames.

Response *Response Status* **U**

REJECT. Doing what the commenter suggests (using an Ethertype) would impact significantly impact the overhead for IET and decrease throughput. Currently, IET provides no change in link throughput for unpreempted frames and minimzes the impact for preempted frames.

Also, if this information was put into the data field fo a frame, that would change the CRC. There is no demonstration of how to do that without weakening the MTTFPA for the resulting frames. It would also require changes to the MAC as it is the MAC that handles frames. The project objectives do not allow that.

The current draft uphods the architecture by not mixing below the MAC content with above the MAC content.

Cl **99** *SC* **99.3.6** *P* **35** *L* **19** # 99

Hajduczenia, Marek Bright House Network

Comment Type **TR** *Comment Status* **R**

What is "the final mPacket"? Likely, "the mPacket containing the final fragment of a frame"

SuggestedRemedy

Per comment - this term is used without definition ...

Response *Response Status* **U**

REJECT. It doesn't define a term. It is a phrase which clearly says the final mPacket of the frame, i.e. the last mPacket - the frame is over. Since the frame is sent in order, that is as clear as the longer phrase.

Cl **99** *SC* **99.4.8** *P* **46** *L* **38** # 195

Marris, Arthur
Cadence Design Syst

Comment Type **TR** *Comment Status* **A**

"shall meet the delay specified elsewhere in this standard" is not an appropriate way to standardize something.

SuggestedRemedy

Replace "the delay specified elsewhere in this standard" with an actual value.

Response *Response Status* **U**

ACCEPT IN PRINCIPLE. It isn't one specific value. Each speed specifies it. We could say "shall meet the delay specified for a MAC Control, MAC and RS based on the MAC operating speed."