

CI **FM** SC **FM** P1 L1 # i-1
 Hajduczenia, Marek Bright House Network

Comment Type **E** Comment Status **A**

Based on IEEE P802.3by entering sponsor ballot in November 2015, IEEE P802.3bq and IEEE P802.3bp entering sponsor ballot in December 2015, the published timeline for IEEE P802.3bq showing approval in June 2016, and the published timeline for IEEE P802.3bp showing approval in August 2016, it seems likely that that IEEE P802.3by will be the second amendment, IEEE P802.3bq will be the third amendment, and IEEE P802.3bn will be the fifth or sixth amendment to IEEE Std 802.3-2015.

SuggestedRemedy

Please change '(Amendment of IEEE Std 802.3(TM)-2015)' to read 'Amendment of IEEE Std 802.3(TM)-2015 as amended by IEEE Std 802.3bw(TM)-2015, IEEE Std 802.3by(TM)-201X, IEEE Std 802.3bq(TM)-201X, IEEE Std 802.3bp(TM)-201X"
 Keep the list updated as project status changes

Response Response Status **C**

ACCEPT IN PRINCIPLE. Add bw and any of the others that are approved for forwarding to RevCom on 18 March 2016. For the others it isn't clear what order they will complete in and IEEE editorial staff can update the list if needed after approval.

CI **FM** SC **FM** P12 L12 # i-2
 Hajduczenia, Marek Bright House Network

Comment Type **E** Comment Status **A**

Suggest that this text be updated based on: (a) the approval of IEEE Std 802.3bw-2015, the likelihood that IEEE P802.3by will be the second amendment, IEEE P802.3bq will be the third amendment, and IEEE P802.3bp will be the fourth amendment to IEEE Std 802.3-2015; (b) use of the (TM) symbol only on the first instance; and (c) alignment of IEEE P802.3br description with other amendment descriptions

SuggestedRemedy

[1] The following text should be inserted prior to the existing text 'IEEE Std 802.3bn(TM)-201x':

IEEE Std 802.3bw-2015

Amendment 1--This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 96. This amendment adds 100 Mb/s Physical Layer (PHY) specifications and management parameters for operation on a single balanced twisted-pair copper cable.

IEEE Std 802.3by-201x

Amendment 2--This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 105 through Clause 112, Annex 109A, Annex 109B, Annex 110A, Annex 110B, and Annex 110C. This amendment adds MAC parameters, Physical Layers, and management parameters for the transfer of IEEE 802.3 format frames at 25 Gb/s.

IEEE Std 802.3bq-201x

Amendment 3--This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 113 and Annex 113A. This amendment adds new Physical Layers for 25 Gb/s and 40 Gb/s operation over balanced twisted-pair structured cabling systems.

IEEE Std 802.3bp-201x

Amendment 4--This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 97 and 98. This amendment adds point-to-point 1 Gb/s Physical Layer (PHY) specifications and management parameters for operation on a single balanced twisted-pair copper cable in automotive and other applications not utilizing the structured wiring plant.

IEEE Std 802.3bn-201x

Amendment 5--This amendment adds the physical layer specifications and management parameters for symmetric and/or asymmetric operation of up to 10 Gb/s on point-to-multipoint Radio Frequency (RF) distribution plants comprising either amplified or passive coaxial media. It also extends the operation of Ethernet Passive Optical Networks (EPON) protocols, such as Multipoint Control Protocol (MPCP) and Operation Administration and Management (OAM).

[2] Insert "Amendment 6--" before the current descriptive text for IEEE Std 802.3br(TM)-201x

[3] Change the description to read: "This amendment includes changes to IEEE Std 802.3-2015 and adds Clause 99.This amendment adds a MAC Merge sublayer and a MAC Merge Service Interface to support for Interspersing Express Traffic over a single Ethernet link."

Response Response Status **C**

ACCEPT IN PRINCIPLE. Add bw and any of the others that are approved for forwarding to

RevCom on 18 March 2016. For the others it isn't clear what order they will complete in and IEEE editorial staff can update the list if needed after approval.

Cl 0 SC 0 P 0 L 0 # i-92
Remein, Duane Futurewei Technologie

Comment Type TR Comment Status R

I concur with comment #33 against D2.2 "... 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

Implement response from either Comment #33 or #31 against D2.2

Response Response Status U

REJECT. The response to comment #33 during WG ballot of Draft 2.2 still applies: 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.

Cl 0 SC 0 P 0 L 0 # i-91
Remein, Duane Futurewei Technologie

Comment Type TR Comment Status D

I concur with comment #31 against D2.2 "This project has failed to live up to the level of participation that was advertised in the PAR: ..."

SuggestedRemedy

See comment #31 D2.2

Proposed Response Response Status Z

REJECT.

This comment was WITHDRAWN by the commenter.

Cl 0 SC 0 P 0 L 0 # i-90
Remein, Duane Futurewei Technologie

Comment Type TR Comment Status R

The terminology in the amendment does not match the agreed objectives for the project. See comment #13 against Draft 2.2 for additional details.

SuggestedRemedy

Update the terminology globally in the draft per the agreed objectives. See comment #13 against Draft 2.2 for details.

Response Response Status U

REJECT. The response to comment #13 during WG Ballot of Draft 2.2 still applies: The main complaint about the initial 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.

Cl 1 SC 1.4.197a P 17 L 8 # i-3
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

"handles express frames." - we define "express traffic" and not "express frames"

SuggestedRemedy

Change to "handles express traffic"

Response Response Status C

ACCEPT.

Cl 1 SC 1.4.339a P 17 L 24 # i-4
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

"handles preemptable frames." - we define "preemptable traffic" and not "preemptable frames"

SuggestedRemedy

Change to "handles preemptable traffic"

Response Response Status C

ACCEPT.

Cl 30 SC 30.2.2.1 P 18 L 20 # i-5
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A contain

This makes very little sense: "If oMACMergeEntity is implemented, the oMACEntity for the express MAC (eMAC) contains an instance of oMACMergeEntity and the oMACMergeEntity contains an instance of oMACEntity for the preemptable MAC (pMAC)." and does not correspond to layering draeing in Clause 99, where two instances (1:many) of MAC are connected to a single MAC Merge connected to a PHY

SuggestedRemedy

Change to: "If oMACMergeEntity is implemented, the oMACEntity for the express MAC (eMAC) and the oMACEntity for the preemptable MAC (pMAC) are connected to an instance of oMACMergeEntity."

Update Figure 30-3, showing many:1 relationship between oMACEntity and oMACMergeEntity, and then creating a 1:many relationship between oMACMergeEntity and oPHYEntity. Remove the secondary pMAC oMACEntity connected to oMACMergeEntity.

Update definition of oMACMergeEntity to read: "If implemented, a single instance of oMACMergeEntity is associated with eMAC and pMAC oMACEntity (see Clause 99). oMACMergeEntity managed object class provides the management controls necessary for the MAC Merge sublayer."

Response Response Status C

ACCEPT.

Cl 30 SC 30.2.5 P 19 L 1 # i-69
 Law, David Hewlett Packard Enter

Comment Type TR Comment Status A contain

*** Comment submitted with the file 88713900003-IEEE_P802d3br_Clause_30_050216.fm attached ***

Taking a top down view, starting at a switch port that supports IET, there would be two subordinate instances of oMACEntity, one would provide management of the pMAC instance and one would provide the management of the eMAC instance. This would be a 'one-to-many relationship' in terms of Figure 30-3, however this is beyond the scope of IEEE 802.3 as the switch port is a MAC Client.

Moving to what is in scope, these two instances of oMACEntity would have a single subordinate instance of oMACMergeEntity to provide management of the single instance of the MAC Merge sublayer for a pMAC and eMAC. This would therefore be a 'many-to-one relationship' in terms of Figure 30-3.

Subordinate to that would be then be oPHYEntity. This is a single instance, except for the case of the MII, which as discussed during the reflector conversation supports the ability to bus PHYs. In that case there would be multiple instances of oPHYEntity, hence this needs to be a 'one to many relationship' in terms of Figure 30-3.

Based on the above I don't think the current IEEE P802.3br draft D3.0 Figure 30-3 reflects this. Starting again at oMACEntity, according to the current figure 30-3, there is a 'one-to-one relationship' to a subordinate instance of oMACMergeEntity which I don't think is correct. Further, subordinate to oMACMergeEntity is a another instance of oMACEntity which to me seems to be circular, since oMACMergeEntity is shown as subordinate to oMACEntity.

SuggestedRemedy

I propose the attached as a replacement. As you will see there is a 'many-to-one relationship' from oMACEntity to oMACMergeEntity and a 'one-to-many relationship' from oMACMergeEntity to oPHYEntity for the case when IET is implemented. When it is not implemented, we have the EPON case of a 'many-to-one relationship' from oMACEntity to oMPEmulation. If not we have the 'one-to-many relationship' from oMACEntity to oPHYEntity. As IET isn't supported by EPON there isn't a need for a relationship from oMACMergeEntity to oMPEmulation.

Response Response Status C

ACCEPT.

Cl 30 SC 30.12.1.1.1 P 22 L 3 # i-6
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status R
Bit numbers should be written in numeric, not in words - it is simpler to read

SuggestedRemedy
change "first" to "1", "two" to "2", etc.

Response Response Status C
REJECT. All the Bit Strings in Clause 30 use ordinal numbers to describe the content of the string. (E.g. 30.3.6.1.6 aOAMLocalConfiguration). Ordinal numbers are unambiguous - there is only one interpretation of first bit. If cardinal numbers were used, one would have to deal with the ambiguity of whether the first bit is bit 0 or bit 1 and it would be inconsistent with the rest of Clause 30.

Cl 30 SC 30.12.2.1.37 P 22 L 52 # i-9
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status R
"A 2-bit integer value used to indicate .." - is there any need for "2-bit" qualifier?

SuggestedRemedy
Remove "2-bit" from the statement, there is no need for it. Same for 30.12.3.1.31, 30.14.1.7,

Response Response Status C
REJECT. When the syntax of the object is INTEGER and the size of the integer is constrained, it is usual to include either the number of bits in the integer or the range of the integer in the Behavior description. See, for example, 30.3.6.1.8 aOAMLocalPDUCOnfiguration and 30.3.6.1.34 aOAMLocalErrSymPeriodConfig

This parameter is only allowed to have a value up to 3 and constraining it to 2-bits does this.

Cl 30 SC 30.12.2.1.37 P 23 L 2 # i-7
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
"I.e., the" - "i.e.," is not needed

SuggestedRemedy
Change to "The"

Response Response Status C
ACCEPT.

Cl 30 SC 30.12.2.1.37 P 23 L 3 # i-8
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
Wrong multiplication sign

SuggestedRemedy
Change "x" to proper multiplication symbol

Response Response Status C
ACCEPT.

Cl 30 SC 30.12.3.1.31 P 23 L 50 # i-62
Gunther, Craig HARMAN INTERNATI

Comment Type E Comment Status A
Add an example of how this attribute is used similar to what was done for aLldpXdot3LocAddFragSize on page 23, lines 2-3.

SuggestedRemedy
Add example

Response Response Status C
ACCEPT IN PRINCIPLE. Add:
the minimum non-final fragment size is
(aLldpXdot3LocAddFragSize + 1) x 64 octets.;

Cl 30 SC 30.14.1.2 P 24 L 33 # i-63
Gunther, Craig HARMAN INTERNATI

Comment Type E Comment Status A
Typo: preemtion

SuggestedRemedy
Replace with "preemption"

Response Response Status C
ACCEPT.

Cl 30 SC 30.14.1.8 P 25 L 53 # i-10
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A

Do not allow for Figure XX-XX to break across the lines.

SuggestedRemedy

Per comment

Response Response Status C

ACCEPT IN PRINCIPLE. It should be a non-breaking space

Cl 30 SC 30.14.1.13 P 26 L 51 # i-11
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status R

Force 30.14.1.13 to move to next page so that the heading is not left stranded

SuggestedRemedy

Per comment

Response Response Status C

REJECT. ATTRIBUTE is not a heading. It is part of the managed object definition. The published standard has page breaks after ATTRIBUTE

See 30.5.1.1.21 aSNROpMarginChnlC, 30.7.1.1.10 aAggPartnerSystemID, and 30.8.1.1.20 aPathSESS for example.

Cl 30 SC 30.14.1.13 P 27 L 5 # i-97
Law, David Hewlett Packard Enter

Comment Type T Comment Status A

Since subclause 99.2.2.1.2 'When generated' states that 'The generation of this primitive is out of scope of this standard.' it is possible that the MM_CTL.request primitive could be generated with the hold_req parameter set to the value HOLD multiple times while requesting preemption. I suspect however that we would only want this counter to increment when there was a transition from hold_req set to RELEASE to hold_req set to HOLD.

SuggestedRemedy

To be clear with our intent with this counter, and since we already map the hold_req parameter of the MM_CTL.request primitive to the variable 'hold' in subclause 99.4.7.3, suggest that the 'BEHAVIOUR DEFINED AS' text be changed to read:

'A count of times of the number of time the variable "hold" (see 99.4.7.3) transitions from FALSE to TRUE.;

Response Response Status C

ACCEPT IN PRINCIPLE.

A count of the number of time the variable hold (see 99.4.7.3) transitions from FALSE to TRUE.;

Cl 30 SC 30.14.1.13 P 27 L 6 # i-98
Law, David Hewlett Packard Enter

Comment Type T Comment Status A

It isn't MM_CTL.request that can have the value HOLD and RELEASE but instead the hold_req parameter supplied by the MM_CTL.request primitive. As stated in IEEE P802.1Qbu changes to subclause 6.7.1, it is when '... a MM_CTL.request(hold_req) primitive is issued to the underlying 802.3 MAC, with a hold_req parameter value of HOLD ...'

SuggestedRemedy

Suggest that:

- [1] The text 'A count of times MM_CTL.request(HOLD) primitive ...' on page 27, line 6 be changed to read 'A count of times the MM_CTL.request primitive is received with a hold_req parameter value of HOLD ...'.
- [2] The text '... when a MM_CTL.request(HOLD) is received ...' on page 43, line 38 be changed to read '... when a MM_CTL.request primitive is received with a hold_req parameter value of HOLD ...'.
- [3] The text '... when MM_CTL.request is received with a value of HOLD and FALSE when MM_CTL.request is received with the value RELEASE.' on page 46, line 29 is changed to read '... when MM_CTL.request is received with a hold_req parameter value of HOLD and FALSE when MM_CTL.request is received with a hold_req parameter value of RELEASE.'
- [4] The text '... MAC Client sends MM_CTL.request(HOLD).' on page 52, line 47 be changed to read '... MAC Client sends MM_CTL.request with a hold_req parameter value of HOLD.'
- [5] The text '... held by an MM_CTL.request(HOLD) sent by ...' on page 52, line 52 be changed to read '... held by an MM_CTL.request with a hold_req parameter value of HOLD sent by ...'.
- [6] The text '... not held by an MM_CTL.request(HOLD) sent by ...' on page 52, line 37 be changed to read '... not held by an MM_CTL.request with a hold_req parameter value of HOLD sent by ...'.
- [7] The text '... is released by an MM_CTL.request(RELEASE) sent ...' on page 53, line 42 be changed to read '... is released by an MM_CTL.request with a hold_req parameter value of RELEASE sent ...'.

Response Response Status C

ACCEPT IN PRINCIPLE.

Make the changes in the suggested remedy except for [1], make the change from I-97 instead:
"A count of the number of times the variable hold (see 99.4.7.3) transitions from FALSE to TRUE.;"

Cl 79 SC 79.3 P 28 L 14 # i-64
Gunther, Craig HARMAN INTERNATI

Comment Type T Comment Status A

Replace TBD subtype with assigned value.

SuggestedRemedy

Assign value to TBD subtype. Also update line 30 on this page. Then remove Editor's Note on lines 16-18.

Response Response Status C

ACCEPT IN PRINCIPLE. The value is 7. (Editor has license to change the value if instructed by the IEEE 802.3 Chief Editor)

Cl 79 SC 79.3 P 28 L 15 # i-12
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Change TBD to actual value

SuggestedRemedy

Per comment + update Figure 79-8 accordingly !
Remove editorial note in line 16-18

Response Response Status C

ACCEPT IN PRINCIPLE. The value should be 7. (Editor has license to change the value if instructed by the IEEE 802.3 Chief Editor)

Cl 79 SC 79.3.7 P 28 L 35 # i-13
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status R

Figure Figure 79-8 is not consistent in the use of "=" symbol. Note that 802.3 OUI field does not use = and others do.

SuggestedRemedy

Add "=" to 802.3 OUI field to separate the value from the field name

Response Response Status C

REJECT. All the existing TLVs in 802.3 also omit the = symbol for the OUI field.

Editor will check on whether there is a reason the "=" should be omitted and if there isn't, will add it.

(Editor's note: The advice of the chief editor was to keep this consistent with the rest of Clause 79 and not make the change.)

Cl 79 SC 79.3.7.1 P 28 L 45 # i-89
Remein, Duane Futurewei Technologie

Comment Type E Comment Status A
Wording: "if more octets are received that were defined"

SuggestedRemedy
Change "that" to "then" so the statement reads: "if more octets are received than were defined"

Response Response Status C
ACCEPT IN PRINCIPLE. "that" to "than"

Cl 79 SC 79.3.7.2 P 29 L 15 # i-14
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
No need for "." at the end of the sentence

SuggestedRemedy
Per comment

Response Response Status C
ACCEPT.

Cl 90 SC 90.4.3.1.1 P 32 L 22 # i-15
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status R
It seems that MM is optional when there is no MAC Merge Layer available. This is not reflected in the text right now.

SuggestedRemedy
Add the statement "The MM parameter is optional." at the beginning of line 22. Remove the statement "The MM parameter is not provided when MAC Merge sublayer is not instantiated." - one could chose to implement it when MAC Merge is not available as well. The same change in 90.4.3.2.1

Response Response Status C
REJECT.
The parameter is not optional. It is required to be used when MAC Merge is instantiated and not used when MAC Merge is not instantiated.

It can not be validly used when MAC Merge is not instantiated as there is no pMAC and eMAC and therefore no value to put in it that makes any sense.

Cl 90 SC 90.5.1 P 33 L 12 # i-16
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
Fix the missing commas around "i.e." - change to ", i.e.,"

SuggestedRemedy
Per comment for 90.5.1 and 90.5.2, including legacy text from 802.3-2015

Response Response Status C
ACCEPT.

CI 99 SC 99.1 P 35 L 15 # i-70
RAN, ADEE Intel Corporation

Comment Type T Comment Status A

The content of the first two sentences is complicated (and seems recursive) - MM sublayer supports something which is achieved by using a MM sublayer.

Also, there are three PLS service interfaces in the figure 99-2; the MM does not simply attach two MACs to a single interface - it merges the traffic, which is quite different.

Also, defining the operation as merging to a PLS service interface would enable cascading MAC Merge sublayers (two MM's merged by a third MM to the RS) - which doesn't seem to be the intent. The merging should only occur once, on the RS.

SuggestedRemedy

Change

"The MAC Merge sublayer supports interspersing express traffic with preemptable traffic. This is achieved by merging the Physical Signaling Sublayer (PLS) service interfaces of an express Media Access Control (MAC) and a preemptable MAC to a single PLS service interface."

to

"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 Reconciliation Sublayer (RS)."

Response Response Status C

ACCEPT IN PRINCIPLE. The "from" text in the suggested remedy doesn't appear in the draft. The text in the draft is:

"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 difference between this and the suggested remedy is replacing the end of the sentence, "a single Physical Signaling Sublayer (PLS) service." with "a single Reconciliation Sublayer (RS)." Make this change.

CI 99 SC 99.1 P 35 L 24 # i-17
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

A primitive does not stop or resume transmission; a primitive causes the transmission of preemptable traffic to be stopped or resumed

SuggestedRemedy

Change "a primitive that either stops or resumes transmission of preemptable traffic, minimizing the latency for express traffic" to a primitive that causes the transmission of preemptable traffic to be either stopped or resumed, minimizing the latency for express traffic"

Response Response Status C

ACCEPT IN PRINCIPLE.

"a primitive that causes the transmission of preemptable traffic to be either held or released, minimizing the latency for express traffic"

because the primitive doesn't unilaterally cause resumption. Express traffic may prevent resumption or there may be no traffic to resume. (See i-18 - "release" will be used for allowing preemptable traffic to flow and "resume" for resuming transmission of a preempted packet.)

CI 99 SC 99.1 P 35 L 26 # i-18
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A resume

Terminology: hold, stop, suspend is used in this clause to mean the very same thing - stop transmission of preemptable traffic until express traffic is done transmitting

SuggestedRemedy

Change "can be held" to "can be suspended". Change is global for Clause 102, including changing the primitive hold_req values to RELEASE and SUSPEND, which are more meaningful in this context

Response Response Status C

ACCEPT IN PRINCIPLE. There are no instances of suspend.

Will replace "stop" with "hold" or "held" as appropriate tense in text. Stop was felt by some to have the implication of aborting transmission by some. Also, HOLD is already used for the primitive in IEEE 802.1Qbu which has completed sponsor ballot.

"Release" will be used for allowing preemptable traffic to flow - so for talking about the effect of the primitive. Release is appropriate because it indicates allowing traffic but there might not be any traffic to send when it occurs.

"Resume" will be used as currently in the state machine for resuming transmission of a preempted packet.

CI 99 SC 99.1 P 35 L 29 # i-99
Law, David Hewlett Packard Enter

Comment Type E Comment Status A

It appears that there are 27 occurrences of 'the preemption capability' and 16 occurrences of 'preemption capability' without the use of a 'the'. As an example page 35 line 29 reads 'When preemption capability ...' yet page 39 line 39 reads 'When the preemption capability ...'.

SuggestedRemedy

Suggest all occurrence of 'preemption capability' should use a 'the', for example the text 'When preemption capability ...' on page 35 line 29 should read 'When the preemption capability ...'.

Response Response Status C

ACCEPT IN PRINCIPLE.

Agree that the should be used in sentences. For table 79-7a function column; and 79-9 and 79-10 TLV variable which are names of the functions of TLV bits, "the" seems unneeded and strange. Suggest leaving "the" out in those instances.

CI 99 SC 99.1 P 35 L 30 # i-100
Law, David Hewlett Packard Enter

Comment Type E Comment Status A

There is only one service primitive defined for the MMSI.

SuggestedRemedy

Suggest that the text '... or the MMSI service primitives to ...' be changed to read '... or the MMSI service primitive to ...'.

Response Response Status C

ACCEPT.

CI 99 SC 99.1 P 35 L 32 # i-20
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Unnecessarily complex dscription: When preemption capability is inactive, the MAC Merge sublayer does not preempt transmission of preemptable packet even if express traffic becomes available. If the MAC Merge sublayer is idle (i.e. at least an interpacket gap has elapsed since the end of transmission of a prior packet) and an express packet becomes available, the MAC Merge sublayer transmits the express packet. Otherwise, the MAC Merge sublayer transmits any presented preemptable packets.

SuggestedRemedy

Change to: When preemption capability is inactive, the MAC Merge sublayer performs multiplexing of data presented by pMAC and eMAC, without suspending transmission of preemptable traffic when express traffic becomes available.

Response Response Status C

ACCEPT IN PRINCIPLE. The suggested text isn't accurate because the MAC Merge sublayer does hold (suspend) preemptable traffic when preemption capability is inactive. For example, if the primitive is used to hold traffic, MAC Merge will not transmit preemptable traffic regardless of whether express traffic is present.

"When preemption capability is active, the MAC Merge sublayer allows frames provided over the express MAC service interface (express traffic) or the MMSI service primitives to interrupt transmission of a preemptable frame provided over the preemptable MAC service interface (preemptable traffic).

When preemption capability is inactive, the MAC Merge sublayer does not allow frames provided over the express MAC service interface (express traffic) or the MMSI service primitives to interrupt transmission of a preemptable frame provided over the preemptable MAC service interface (preemptable traffic).

Regardless of whether preemption capability is active, the MAC Merge sublayer allows frames provided over the express MAC service interface (express traffic) or the MMSI service primitives to prevent start of transmission of a preemptable frame provided over the preemptable MAC service interface (preemptable traffic)."

[Editor's note: after doing some clean-up to use "express traffic" in place of the longer phrase after the first occurrence of "frames provided over the express MAC service interface (express traffic)" and applying changes from other comments, the text used is: "When the preemption capability is active, the MAC Merge sublayer allows frames provided over the express MAC service interface (express traffic) or the MMSI service primitive to interrupt transmission of a frame provided over the preemptable MAC service interface.

When the preemption capability is inactive, the MAC Merge sublayer does not allow express traffic or the MMSI service primitive to interrupt transmission of a frame provided over the preemptable MAC service interface.

Regardless of whether the preemption capability is active, the MAC Merge sublayer allows express traffic or the MMSI service primitive to prevent the start of transmission of frames provided over the preemptable MAC service interface (preemptable traffic)."]

CI 99 SC 99.1 P 35 L 34 # i-19
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A
Mixing "preemptable packet" and "preemptable traffic" - stick with "preemptable traffic", which is what is more generic and defined

SuggestedRemedy

Per comment

Response Response Status C

ACCEPT IN PRINCIPLE. Fixed by the changes in i-20.

"When the preemption capability is active, the MAC Merge sublayer allows frames provided over the express MAC service interface (express traffic) or the MMSI service primitive to interrupt transmission of a frame provided over the preemptable MAC service interface.

When the preemption capability is inactive, the MAC Merge sublayer does not allow express traffic or the MMSI service primitive to interrupt transmission of a frame provided over the preemptable MAC service interface.

Regardless of whether the preemption capability is active, the MAC Merge sublayer allows express traffic or the MMSI service primitive to prevent the start of transmission of frames

CI 99 SC 99.1 P 35 L 34 # i-71
RAN, ADEE Intel Corporation

Comment Type E Comment Status A
"i.e." should be followed by a comma.

SuggestedRemedy

per comment.

Response Response Status C

ACCEPT IN PRINCIPLE. This instance of i.e. was removed in response to i-20.

CI 99 SC 99.1 P 35 L 38 # i-93
Remein, Duane Futurewei Technologie

Comment Type TR Comment Status R

It has been observed by others that options tend to become requirements in the market. There are numerous RFPs that require Ethernet features that are optional just because the option appears in the standard and it is easier to require all the bells and whistles than to pick and choose, especially if there is a slight chance that the feature will be needed "someday". This is overriding fear with this project; that it will become a required feature for all MACs creating in effect a Tax on Ethernet. It should be made clear that this feature should not be required of MACs not intended on the targeted application (automotive and similar applications).

SuggestedRemedy

Append to the para starting "Preemption capability is most useful at lower operating speeds" the following:
"Therefore, Express Traffic features should not be implemented in very high speed MACs (e.g., at rate greater than 5 gaps). Furthermore Express Traffic can place a burden on lower speed MACs that do not need the advantages of the interspersed express traffic feature should only be included in MACs targeting applications (such as automotive and industrial) that receive significant benefits from this feature.

Response Response Status U

REJECT.
Automotive and industrial were two markets that justified starting the work, but IEEE P802.1Qbu Preemption and IEEE P802.3br IET are useful for other markets as well. Other examples including pro-audio and video, building automation, smart grid, power generation and front haul networks.

The front haul network use case requires low enough latency and high bandwidth such that preemption provides for longer reach needs at 10 Gb/s. See: <http://www.ieee802.org/1/files/public/docs2016/cm-farkas-profiles-A-and-B-0316-v01.pdf>

CI 99 SC 99.1 P 35 L 39 # i-108
 Scruton, Peter

Comment Type T Comment Status R

Calculation of duration of a 2000 byte frame should include the preamble and SFD (additional 64 bit times). So this would make the delay at 100 be 160.64 and for 1000 would be 16.064 uS.

SuggestedRemedy

"For example, the duration of a 2000 octet packet (including Preamble and SFD) on a 100 Mb/s link is 160.64 us and on a 1 Gb/s link is 16.064 us."

Response Response Status C

REJECT. The packet includes the preamble and SFD so this is the transmission time of a 2000 octet packet.

Also as far as an upper bound on the time saved by preemption, 160 and 16 us are the upper bound since preemption can't occur until 64 octets after the SFD is sent so the preamble and SFD don't add to the additional delay.

CI 99 SC 99.1 P 35 L 40 # i-109
 Scruton, Peter

Comment Type T Comment Status R

I don't see this as the upperbound as the upperbound should include the preamble and SFD and the time to transmit the IPG as that would be the extra delay if they were both presented at the same time. Which would be 16000 + 64 + 96 = 16160 or 161.6uS at 100.

SuggestedRemedy

Modify sentence to: "This, along with the time associated for an IPG, provides an ... "

Response Response Status C

REJECT. This is an upper bound on "the additional delay before a MAC Client can send an Express frame when preemption capability is not used." A packet can't be preempted until 64 octets after the SFD. Whether transmission of a packet ends or the packet is preempted, there will be an IPG followed by the preamble and SFD of the express packet. So the IPG, preamble and SFD are not additional delay. That delay occurs regardless of whether preemption capability is used or not.

CI 99 SC 99.1 P 35 L 40 # i-21
 Hajduczenia, Marek Bright House Network

Comment Type ER Comment Status A

Please use proper symbols for "us"

SuggestedRemedy

Per comment

Response Response Status C

ACCEPT.

CI 99 SC 99.1 P 35 L 40 # i-22
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Unclear what "This is an upper bound" is intended to mean - the time it takes 2k frame to pass through MAC, 16 us, or 160us.

SuggestedRemedy

Please clarify what "this" means in this context, at best listing it once more

Response Response Status C

ACCEPT IN PRINCIPLE. "The time to transmit a maximum length packet is an upper bound..."

CI 99 SC 99.1 P 35 L 53 # i-23
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A

Extra "." in "Frame Preemption. and IEEE Std 802.1Qbv"

SuggestedRemedy

Change to "Frame Preemption, and IEEE Std 802.1Qbv"

Response Response Status C

ACCEPT IN PRINCIPLE. Since it is A and B, no comma is needed either.

Cl 99 SC 99.1 P 36 L 9 # i-73
RAN, ADEE Intel Corporation

Comment Type E Comment Status A

MAC is expanded as "MAC -- MEDIA ACCESS CONTROL" twice. The term appears in upper boxes unexpanded.

Other acronyms are expanded in the legend.

SuggestedRemedy

Change "MAC -- MEDIA ACCESS CONTROL" to "MAC" within the figure (twice) and add "MAC = MEDIA ACCESS CONTROL" to the legend.

Response Response Status C

ACCEPT.

Cl 99 SC 99.1 P 36 L 16 # i-24
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

The separation of PHY into PCS, PMA, and PCS is not applicable to all PHY types defined in 802.3 - some do not separate PMA and PMD, combining them into a single PHY instead.

SuggestedRemedy

To be more generic, we should show really a single box under xMII and call it PHY instead. The presence of PCS, PMA, PMD, and any potential extenders between sublayers is irrelevant for .3br.

Response Response Status C

ACCEPT.

Cl 99 SC 99.1 P 36 L 18 # i-72
RAN, ADEE Intel Corporation

Comment Type TR Comment Status A

Figure 99-1 includes a PMD sublayer, but a PMD is not defined in some PHYs (especially those used in automotive environments, defined in 802.3bp and 802.3bw).

SuggestedRemedy

Delete the PMD sublayer from the figure and the legend.

Consider deleting the division of PHY to sublayers and changing the box content to "PHY sublayers".

Response Response Status C

ACCEPT. Delete the division of the PHY into sublayers

Cl 99 SC 99.1 P 36 L 32 # i-74
RAN, ADEE Intel Corporation

Comment Type TR Comment Status R

"A MAC Control Sublayer that is the client of an eMAC or a pMAC shall not generate PAUSE"

This is stated as a normative requirement on another sublayer. That breaks the layer separation. What happens if someone builds a system with multiple IPs and the MAC control does generate PAUSE? is the behavior undefined?

The normative statement should apply to the MAC Merge sublayer behavior.

Also, the PICS item for that state ment (MM14) is stated as "No PAUSE | 99.1 | MAC Control sublayer shall not generate PAUSE". This is too broad.

Is there any concern of PAUSE (in either direction) when the device has to "interoperate with devices that do not implement the MAC Merge sublayer"?

SuggestedRemedy

Change

"A MAC Control Sublayer that is the client of an eMAC or a pMAC shall not generate PAUSE"

to (something along the lines of)

"The MAC merge sublayer does not support PAUSE request from either the pMAC and the eMAC. Any PAUSE request received from a MAC Control Sublayer that is a client of the eMAC or the pMAC shall not be placed on the PLS service interface".

Alternatively (or additionally), bring in Annex 31B and add a qualification that PAUSE is disabled when MAC merge is used (similar to the current last paragraph of 31B.1).

If the normative statement stays, reword it to refer to the behavior of the MAC merge sublayer.

Response Response Status C

REJECT. We have other requirements that restrict the applicability of using one sublayer with another.

For example, when implementing IET, the MACs must be full duplex.

The requirement stated here is a requirement on implementing IET - which involves multiple sublayers. It is not a requirement on the MAC Merge sublayer.

MAC Merge shouldn't be parsing the contents of MAC packets to decide to stop some from being sent. That would be a layer violation.

PAUSE won't work with IET because PAUSE operation requires that the PAUSE frame is received on the MAC Control sublayer that is transmitting the traffic to be PAUSED and a PAUSE sent by the preemptable MAC's MAC Control Sublayer would be received on the express MAC MAC Control sublayer when the sending MAC has preemption capability disabled.

PAUSE is already only applicable when both ends of the link support PAUSE and normally is enabled after auto-negotiation determines that both ends of the link support it. PAUSE only works if both sides of the link support it. IET doesn't change this.

Cl 99 SC 99.1 P 36 L 32 # i-25
Hajduczenia, Marek Bright House Network

Comment Type ER Comment Status A
Sentence could use wording improvement

SuggestedRemedy

"When attached to an eMAC or a pMAC, the MAC Control Sublayer shall not generate PAUSE" - this emphases where the DUT is and what it does
Update PICs

Response Response Status C
ACCEPT.

Cl 99 SC 99.1 P 36 L 35 # i-26
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status R

It is actually broader than that: Devices that implement the MAC Merge sublayer interoperate with devices that do not implement the MAC Merge sublayer.

SuggestedRemedy

Change to: Devices that implement and enable the MAC Merge sublayer interoperate with both devices that implement but disable the MAC Merge sublayer, and devices that do not implement the MAC Merge sublayer at all.

Response Response Status C

REJECT. We don't disable the MAC Merge sublayer. We disable or enable preemption capability.

Someone might make a device that can be configured to instantiate MAC Merge and the second MAC or not, but that is outside the standard.

The point here is that this is backwards compatible.

Cl 99 SC 99.1.2 P 38 L 1 # i-27
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status R

There is no reason to separate receive direction into "EXPRESS FILTER" and "RECEIVE PROCESSING" blocks - multiple SDs can run inside of a single function block, with no issues at all

SuggestedRemedy

Merge "Express Filter" and "Receive Processing" into a single block "Receive Processing" and source all PLS_DATA.indication, PLS_DATA_VALID.indication, and PLS_SIGNAL.indication signals for pMAC and eMAC from there.
Align description accordingly

Response Response Status U

REJECT. Either way would be valid. It is a matter of what blocks to break things into for easier consumption by the reader and breaking apart the two blocks of receive functionality aids in that.

Cl 99 SC 99.2 P 38 L 44 # i-101
Law, David Hewlett Packard Enter

Comment Type E Comment Status A

Not sure what providing a service 'on the MMSI' means.

SuggestedRemedy

Suggest that '... specifies the services provided on the MMSI by the MAC Merge ...' be either changed to read '... specifies the services provided across the MMSI by the MAC Merge ...' or preferable '... specifies the services provided by the MAC Merge ...'.

Response Response Status C

ACCEPT IN PRINCIPLE.
'... specifies the services provided by the MAC Merge ...'.

CI 99 SC 99.2.1 P 38 L 46 # i-75
 RAN, ADEE Intel Corporation

Comment Type T Comment Status A

The hierarchy structure of the service interface specification is unnecessarily complex.

This subclause is titled "MMSI" which is the same as its parent subclause 99.2. Additionally, 99.2 body has only one sentence, which seems to be repeated in 99.2.1 (in a weird phrasing - the MMSI doesn't specify, the subclause does).

The last 3 paragraphs of 99.2.1 seem to be unnecessary as well, they practically just point to 99.2.2 that immediately follow.

It seems that cleaning 99.2.1 and merging it with its parent, and making 99.2.2 shallower, would help readers. That would also do away with the "shall" that really goes without saying (service interfaces are abstract so the "shall" doesn't mean anything).

SuggestedRemedy

Delete the first sentence of 99.2.1 and the last 3 paragraphs (starting from "The following primitive is defined" and including the "shall"). Merge the rest of 99.1.1 into 99.1.

Delete current 99.2.2 and promote 99.2.2.1 to become 99.2.1.

The hierarchy would become
 99.2 MAC Merge Service Interface (MMSI)
 99.2.1 MM_CTL.Request
 99.2.1.1 Semantics
 (...)

Delete PICS item MM1.

Response Response Status C

ACCEPT IN PRINCIPLE. All of the service interface specifications in 802.3 that the editor checked have a list of the primitives for the interface before starting the definition of each service primitive.

Delete the heading 99.2.1 MMSI and the first sentence of 99.2.1 and the last sentence of 99.2.1. Merge the rest of 99.2.1 into 99.2.

Delete 99.2.2 and promote 99.2.2.1 to 99.2.1

Delete PICS item MM1

CI 99 SC 99.2.2.1 P 39 L 9 # i-102
 Law, David Hewlett Packard Enter

Comment Type TR Comment Status A resume

Subclause 99.2.2.1 'MM_CTL.request' states that the primitive defines a request from a MAC Client 'to stop or release transmission of preemptable traffic.'. Subclause 99.2.2.1.1 'Semantics' states that the value of hold_req=HOLD 'causes the MAC Merge sublayer to stop transmission of preemptable traffic.'

In both cases isn't it more granular, that is setting hold_req=HOLD won't stop the transmission of preemptable traffic, specifically a mPackets containing preemptable traffic being transmitted when hold_req=HOLD will be completed, instead it will prevent the initiation of transmission of any further mPackets containing preemptable traffic.

Based on this, rather than talking about stopping and starting the transmission of preemptable traffic, wouldn't it be more accurate to talk in terms of stopping and starting in initiation of transmission of mPackets containing preemptable traffic.

Similarly 99.2.2.1.3 'Effect of receipt' states that HOLD 'prevents starting transmission of pMAC packets until this primitive is received with the value RELEASE.'. Again, isn't it the transmission of mPackets by the MAC merge sublayer containing preemptable traffic, rather than packets by the pMAC instance, that are being controlled.

Finally suggest that either 'release' or 'resume' be used consistently in respect to restarting transmission, suggest that 'resume' be used.

SuggestedRemedy

Suggest that:

[1] The text '... to stop or release transmission of preemptable traffic.' on page 39 line 9 be changed to read '... to stop or resume the initiation of transmission of mPackets (see 99.3.1) containing preemptable traffic.'

[2] The text '... to stop transmission of preemptable traffic.' on page 39 line 19 be changed to read '... to stop the initiation of transmission of mPackets containing preemptable traffic.'

[3] The text '... to resume transmission of preemptable traffic.' on page 39 line 21 be changed to read '... to resume the initiation of transmission of mPackets (see 99.3.1) containing preemptable traffic.'

[4] The text '... allow preemption and prevents starting transmission of pMAC packets until this primitive is received with the value RELEASE.' on page 39 line 31 be changed to read '... allow preemption and prevents initiation of transmission of mPackets containing preemptable traffic until this primitive is received with the value RELEASE.'

Response Response Status C

ACCEPT IN PRINCIPLE. This text has been clarified by a number of other comments. Comment i-18 will replace "stop" with "hold" so the draft uses a consistent word, "hold". "hold transmission of preemptable traffic" includes preempting when preemption is allowed

and not starting new frames.

In 99.2.2.1.3 Effect of receipt just below, there are additional details on how transmission of preemptable traffic is held.:

"Receipt of the primitive with the value HOLD causes preemption if the current conditions allow preemption and prevents starting transmission of pMAC packets until this primitive is received with the value RELEASE." Comment i-65 clarifies this by replacing "if the current conditions" with "when conditions allow" so the resulting text will be:

"Receipt of the primitive with the hold_req parameter set to the value HOLD causes preemption when the conditions allow preemption and prevents starting transmission of pMAC packets."

CI 99 SC 99.2.2.1.1 P 39 L 18 # i-103

Law, David Hewlett Packard Enter

Comment Type E Comment Status A

Suggest that since the first sentence of the paragraph states that the hold_req parameter can take the two values HOLD and RELEASE it isn't necessary to include 'hold_req=' in the text. Also suggest this and the next paragraph be one paragraph.

SuggestedRemedy

Suggest the text:

'The value of hold_req=HOLD causes ... preemptable traffic.

The value of hold_req=RELEASE allows ...'

Be changed to read:

'The value HOLD causes ... preemptable traffic. The value RELEASE allows ...'

Response Response Status C

ACCEPT.

CI 99 SC 99.2.2.1.1 P 39 L 19 # i-28

Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

"stop transmission of preemptable traffic" - looking at associated SDs, it seems more of "suspending" transmission than stopping it - the frame is suspended mid-flight until MAC Merge is done with express traffic

SuggestedRemedy

Change to "suspend transmission of preemptable traffic"

Response Response Status C

ACCEPT IN PRINCIPLE. Will use "hold" rather than "suspend" for consistency. (see i-18)

CI 99 SC 99.2.2.1.3 P 39 L 30 # i-65

Gunther, Craig HARMAN INTERNATI

Comment Type E Comment Status A

Clause 99.4.4 specifies that a frame will not be preempted until at least 60 octets have been sent. What will happen if 10 octets of a 2000 octet frame have been transmitted and a MM_CTL.request(HOLD) is received? The "current conditions" will not allow preemption. I could interpret this subclause as saying that the HOLD will not take effect until the entire 2000 octets have been transmitted because the "current conditions" did not allow it when the HOLD was received. Is that the intended behavior? Or would the HOLD cause preemption to occur as soon as 60 octets of the 2000 octets have been transmitted?

SuggestedRemedy

Discuss and clarify if necessary.

Response Response Status C

ACCEPT IN PRINCIPLE.

As is clear in the state machines, if preemption capability is active, preemption will occur after 60 octets frame have been sent in the current mPacket if at least 64 octets of the frame remain.

Change to "when the conditions allow" so the sentence becomes:

"Receipt of the primitive with the hold_req parameter set to the value HOLD causes preemption when the conditions allow preemption and prevents starting transmission of pMAC packets."

CI 99 SC 99.2.2.1.3 P 39 L 30 # i-104

Law, David Hewlett Packard Enter

Comment Type T Comment Status A

It isn't MM_CTL.request that can have the value HOLD and RELEASE but instead the hold_req parameter supplied by the MM_CTL.request primitive.

SuggestedRemedy

Suggest that:

[1] The text 'Receipt of the primitive with the value HOLD causes ...' should be changed to read 'Receipt of the primitive with the hold_req parameter set to the value HOLD causes

[2] The text 'The receipt of this primitive with the value RELEASE allows ...' should be changed to read 'The receipt of this primitive with the hold_req parameter set to the value RELEASE allows ...'.

Response Response Status C

ACCEPT.

CI 99 SC 99.2.2.1.3 P 39 L 34 # i-29
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Repetition - The receipt of this primitive with the value RELEASE allows MAC Merge sublayer to transmit packets from the pMAC when the eMAC does not have a packet to transmit. This is stated before in different ways already

SuggestedRemedy

Change content of 99.2.2.1.3 to read: "Receipt of the primitive with the value HOLD causes the MAC Merge sublayer to suspend transmission of preemptable traffic. Receipt of the primitive with the value RELEASE causes the MAC Merge sublayer to resume transmission of preemptable traffic, and multiplex preemptable and express traffic as it becomes available."

Response Response Status C

ACCEPT IN PRINCIPLE. There have been several comments relevant to this text (i-65 and i-102). The latter were on making it clear that receipt of the primitive with the parameter value of HOLD can cause preemption and the suggested remedy isn't explicit about that.

Change content of 99.2.2.1.3 to read: "Receipt of the primitive with the hold_req parameter set to the value HOLD causes preemption when the conditions allow preemption and prevents starting transmission of pMAC packets.

Receipt of the primitive with the hold_req parameter set to the value RELEASE allows transmission of packets from the pMAC when the eMAC does not have a packet to transmit."

CI 99 SC 99.3.1 P 40 L 1 # i-76
RAN, ADEE Intel Corporation

Comment Type E Comment Status R

99.3.1 is really a summary of details in 99.3.2 to 99.3.6. It seems that it should be merged into 99.3.

SuggestedRemedy

per comment.

Response Response Status C

REJECT. It is showing the format of the mPacket. That includes the order of the fields which isn't included in 99.3.2 to 99.3.6 so it isn't a summary of them.

This is consistent with 3.1.1 for the MAC Frame format and 79.1.1 for the LLDP Frame Format

CI 99 SC 99.3.1 P 40 L 3 # i-31
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Text is not complete: Figure 99-4(a) shows the format of an mPacket containing an express packet or the initial fragment of a preemptable packet.

SuggestedRemedy

Change to: Figure 99-4(a) shows the format of an mPacket containing an express packet, a complete preemptable packet, or the initial fragment of a preemptable packet.

Response Response Status C

ACCEPT.

CI 99 SC 99.3.1 P 40 L 20 # i-30
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R

Notes under Figure 99-4 are not needed, given the text in lines 2-3 is already in place and describes each case completely

SuggestedRemedy

Strike notes in Figure 99-4

Response Response Status C

REJECT. They aren't notes (notes say "note"). They are captions. The captions act as subtitles for the figure. One could by the same token say that Figure and Table titles are unnecessary because the text that references them says what they are but the titles aid the reader.

CI 99 SC 99.3.1 P 40 L 27 # i-32
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

"The format is indicated by the SMD (see 99.3.3)." - the format of what?

SuggestedRemedy

Change to: "The mPacket format is indicated by the value of SMD (see 99.3.3).

Response Response Status C

ACCEPT.

Cl 99 SC 99.3.2 P 40 L 31 # i-33
 Hajduczenia, Marek Bright House Network
 Comment Type T Comment Status A
 "The preamble ... contains preamble octets." doohoo ?
 SuggestedRemedy
 Change three instances of "preable octets" to just "octets"
 Response Response Status C
 ACCEPT.

Cl 99 SC 99.3.2 P 40 L 33 # i-110
 Scruton, Peter
 Comment Type T Comment Status A
 "0x55 (binary 10101010)." The binary transmission order is 10101010 as IEEE 802.3-2012 subclause 3.1.1 states transmission order of a byte as LSB to MSB, but 0x55 in binary is not 10101010.
 SuggestedRemedy
 change to something like: "0x55 (which would create a bit order transmission of 10101010, normal preamble)."
 Response Response Status C
 ACCEPT IN PRINCIPLE. (transmitted in order from left to right 10101010)

Cl 99 SC 99.3.3 P 40 L 46 # i-34
 Hajduczenia, Marek Bright House Network
 Comment Type ER Comment Status A
 - additional fragment counter octet (frag_count) following - Figure 99-4 shows FRAG_COUNT and not frag_count
 SuggestedRemedy
 Change frag_count to FRAG_COUNT
 Response Response Status C
 ACCEPT IN PRINCIPLE. Figure 99-4 has it in all caps because format for figures in 802.3 such as this use all caps for the labels. See for example Fig 3-1, Fig 3-2 and Fig 79-1. Preamble and octet are also all caps in in the figure as they are in the Clause 3 figures but they are not all caps in text.
 There is one case of FRAG_COUNT in the discription of the rxFragCnt variable (page 47, line 56) that is all caps. Replace with lower case.

Cl 99 SC 99.3.3 P 41 L 1 # i-35
 Hajduczenia, Marek Bright House Network
 Comment Type E Comment Status A
 "SMD-S refers to any of the four SMD values ..." - likely, it is intended to be an introduction of a term
 SuggestedRemedy
 Change to: "The term "SMD-S" designates any of the four SMD values ..."
 The same change for "SMD-C" in the following line
 Response Response Status C
 ACCEPT.

Cl 99 SC 99.3.3 P 41 L 9 # i-36
 Hajduczenia, Marek Bright House Network
 Comment Type T Comment Status A
 The column "Encoding" really shows the "Value" of specific SMD code
 SuggestedRemedy
 Change "Encoding" to "Value"
 The same in Table 99-2
 Response Response Status C
 ACCEPT.

Cl 99 SC 99.3.4 P 41 L 36 # i-37
 Hajduczenia, Marek Bright House Network
 Comment Type ER Comment Status R
 When speaking about frag_count (variable) it should be written in lower case. When speaking about the field, it should be capitalized
 SuggestedRemedy
 Change "The frag_count field" to "The FRAG_COUNT field"
 Response Response Status U
 REJECT. There is no convention in IEEE 802.3 to do that. See for example the frame fields in Clause 3 and in Clause 79. They are upper case in the figures and lower case or initial caps when in text.

Cl 99 SC 99.3.4 P 41 L 44 # i-38
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
Too many values: "The valid values of frag_count values

SuggestedRemedy
Change to: The valid frag_count values ...

Response Response Status C
ACCEPT.

Cl 99 SC 99.3.4 P 42 L 1 # i-39
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status R
Inconsistency capitalization: Frag_count values

SuggestedRemedy
Change to: frag_count values. Also change the capitalization in the column name in Table 99-2

Response Response Status C
REJECT. It is capitalized because Figure titles have an initial cap as do column titles.

Cl 99 SC 99.3.5 P 41 L 49 # i-41
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A
Which CRC is intended here: CRC32 from the original MAC packet, or CRC calculated by MAC Merge?

SuggestedRemedy
Suggest to use CRC in meaning of CRC32 from original MAC frame. Use mCRC to designate value calculated by MAC Merge sublayer
In this line, change CRC to mCRC. Update Figure 99-4 (both a and b).
Update 99.3.6 accordingly:
99.3.6 mCRC

In the final mPacket of a preemptable frame, the mCRC field contains the last 4 octets of the MAC frame (the FCS field, containing the original CRC of the packet). Otherwise, the mCRC field contains a cyclic redundancy check and indication of whether this mPacket is the final fragment of a preemptable frame. The mCRC shall be calculated on the octets of the frame from the first octet of the frame (i.e., the octet following the SFD sent by the pMAC) to the last octet transmitted in that mPacket by:
-- performing steps a) through d) in 3.2.9 and then
-- XORing the calculated 32-bit value with 0x0000 FFFF.
NOTE--0x0000 is XORed with two octets that contain the higher order coefficients of the mCRC and 0xFFFF is XORed with the two octets that contain the lower order coefficients of the mCRC.

Response Response Status C
ACCEPT IN PRINCIPLE. Change this occurrence of "CRC" to "FCS". The CRC of a MAC frame already has a distinct name, FCS, which is what should have been used here.

The draft uses FCS to refer to the CRC computed by the MAC and mCRC for the CRC added when MAC Merge preempts a packet, but the wrong term was used here. The CRC field is called CRC field because it can contain an FCS or an mCRC.

CI 99 SC 99.3.5 P 41 L 49 # i-40
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A
 "The minimum size of the mData field is 60 octets." - what happens if the last fragment is smaller than 60 octets? Does it get padded to meet this requirement? It is not covered right now here

SuggestedRemedy
 Clarify whether padding takes place, or whether it is prohibited altogether.

Response Response Status C
 ACCEPT IN PRINCIPLE.
 There is no padding - see the first paragraph of 9.4.4. The rules for preemption ensure that preemption only occurs when there are at least 60 of the packet data field left to transmit.

The Transmit Processing State Diagram ensures that this requirement is always met without padding through the presence of the MIN_REMAIN function in the preempt variable.

There is no PAD field so it is already clear that no PAD is allowed.

We will add a note that the state diagram ensures that the requirement is met and padding is never used.

CI 99 SC 99.3.6 P 40 L 22 # i-111
 Scruton, Peter

Comment Type T Comment Status A
 "The mCRC shall be calculated on the octets of the frame from the first octet of the frame (i.e., the octet following the SFD sent by the pMAC) to the last octet transmitted in that mPacket by:" I find this potentially ambiguous. It states "The mCRC shall be calculated on the octets of the frame from the first octet of the frame " ... " to the last octet transmitted in that mPacket" ... The last octet of the mPacket is the last octet of the mCRC. The last octet of the frame in that mPacket is 4 bytes earlier.

SuggestedRemedy
 I suggest the following wording to disambiguate: "The mCRC shall be calculated on the octets of the frame from the first octet of the frame (i.e., the octet following the SFD sent by the pMAC) to the last octet transmitted prior to the mCRC field in that mPacket by:"

Response Response Status C
 ACCEPT IN PRINCIPLE. It isn't ambiguous because it says that the mCRC is "calculated on the octets of the frame". The mCRC octets are not octets of the frame.

We will add "of the frame" before "transmitted" to emphasize that.

The mCRC shall be calculated on the octets of the frame from the first octet of the frame (i.e., the octet following the SFD sent by the pMAC) to the last octet of the frame transmitted in that mPacket

CI 99 SC 99.3.6 P 42 L 27 # i-77
 RAN, ADEE Intel Corporation

Comment Type E Comment Status A
 Space in hexadecimal numbers preceded by 0x is a weird convention. The example in 1.2.5 (Hexadecimal notation) does not include a space.

SuggestedRemedy
 Delete the space.

Response Response Status C
 ACCEPT.

CI 99 SC 99.3.6 P 42 L 27 # i-66
 Gunther, Craig HARMAN INTERNATI

Comment Type E Comment Status R
 Is there ever a situation where the mCRC of an intermediate fragment (not the last fragment) could match the CRC of the packet fragments that have been reassembled so far? I'm wondering about a situation where the reassembly algorithm could believe it has seen the final mPacket when it really hasn't. Or vice versa, where the final fragment contains the original CRC that would match the calculated mCRC.

SuggestedRemedy
 If the answer is no, it can never make that mistake, I would suggest it would be worth adding a note about that.

If it ever could occur we may need to introduce another SMD primitive that would be used in those rare cases when an mCRC would match the original CRC. The transmitter would know that and send one more packet with this new SMD that tells the receiver there is no data in the new SMD packet, but it marks the completion of the previous fragmented packet.

Response Response Status C
 REJECT. The answer is no, the two values will never match. The mCRC is calculated by doing a CRC calculation over all the octets of the frame that have been transmitted. The FCS is calculated by doing the same CRC calculation over all the octets of the frame. At that point the results of those two calculations over the same byte string are the same.

If it is an mCRC, 0x0000 FFFF is XORed with that result to produce the mCRC.

If it is an FCS, 0xFFFF FFFF is XORed with the result to produce the FCS

For A not equal to B,
 C XOR A will never equal C XOR B.
 The FCS of a frame never matches the mCRC computation for the bytes of the frame.

The standard isn't a tutorial. There is no need for a note.

Cl 99 SC 99.4 P 42 L 38 # i-78
 RAN, ADEE Intel Corporation
 Comment Type E Comment Status A
 "Any packet"... disagrees with "are received".
 SuggestedRemedy
 Change "are received to "is received".
 Response Response Status C
 ACCEPT.

Cl 99 SC 99.4.1 P 42 L 45 # i-42
 Hajduczenia, Marek Bright House Network
 Comment Type T Comment Status A
 It is not possible for the same packet to be present on pMAC and eMAC: "the MAC Merge sublayer transmits packets rather than mPackets. If both the eMAC and pMAC have a packet ready to transmit"
 SuggestedRemedy
 Change to "the MAC Merge sublayer transmits packets rather than mPackets. If both the eMAC and pMAC have packets ready to transmit
 Response Response Status C
 ACCEPT.

Cl 99 SC 99.4.1 P 42 L 46 # i-43
 Hajduczenia, Marek Bright House Network
 Comment Type T Comment Status R
 What you're trying to describe here is simple first-come-first-serve interleaving: If both the eMAC and pMAC have a packet ready to transmit and no packet is being transmitted, the eMAC packet istransmitted. If a pMAC packet is being transmitted and the eMAC has a packet to transmit, the packet from the eMAC is transmitted after transmission of the pMAC packet completes.
 SuggestedRemedy
 Change to: If both the eMAC and pMAC have packets ready to transmit, the eMAC packets are prioritized over pMAC packets. Otherwise, either MAC instance is allowed to transmit packets.
 Response Response Status C
 REJECT. The current text is accurate. The suggested remedy is inaccurate because it says "Otherwise, either MAC instance is allowed to transmit packets."
 The variable hold might be TRUE and then the pMAC is not allowed to transmit packets.

Cl 99 SC 99.4.3 P 43 L 13 # i-44
 Hajduczenia, Marek Bright House Network
 Comment Type T Comment Status A
 Unclear what this :Verification" is.
 SuggestedRemedy
 Change "Verification checks that the link can support the preemption capability." to "The Verification function (see Figure 99-3) confirms whether the link supports the preemption capability."
 Change all instances of "verification" (when used as noun, referring to the verification function) to "verification function"
 Response Response Status C
 ACCEPT IN PRINCIPLE. Use:
 Verification (see Figure 99-3) confirms whether the link supports the preemption capability.
 Also add a reference to Figure 99-3 in the subclauses for the other blocks in that figure.

Cl 99 SC 99.4.3 P 43 L 15 # i-45
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Explains the low level mechanics, without referring providing higher level explanation of what is going on: If the preemption capability is enabled and has not been verified, MAC Merge sublayer initiates transmission of a verify mPacket.

SuggestedRemedy

Change:
If the preemption capability is enabled and has not been verified, MAC Merge sublayer initiates transmission of a verify mPacket. A verify mPacket has 7 octets of preamble (0x55), an SMD-V, 60 octets of 0x00 and an mCRC. Transmission of a verify packet is repeated if no response is received.
When an mPacket with an SMD-V and a correct mCRC is received, a response mPacket is sent. A response packet has 7 octets of preamble (0x55), an SMD-R, 60 octets of 0x00 and an mCRC.
When an mPacket with an SMD-R and a correct mCRC is received, the preemption capability is verified.
to
If the preemption capability is enabled but has not been verified yes, the MAC Merge sublayer initiates the verification function. The verification function relies on the transmission of verify mPacket and receipt of response mPacket to confirm that the remote station supports the preemption capability. The format of verify mPacket and response mPacket is shown in Figure 99-4(a), with the SMD values defined in Table 99-1.

Fix the name in Table 99-1: change "respond packet" to "response packet" for consistency of naming in 99.4.3

Response Response Status C

ACCEPT IN PRINCIPLE.

"If the preemption capability is enabled but has not been verified yet, the MAC Merge sublayer initiates the verification function. The verification function relies on the transmission of verify mPacket and receipt of respond mPacket to confirm that the remote station supports the preemption capability. The format of verify mPacket and respond mPacket is shown in Figure 99-4(a), with the SMD values defined in Table 99-1 and an mData field containing 60 octets of 0x00."

Search for cases of is "Response" and replace with "Respond" when referring to the packet and state machine because it is more parallel to "Verify". (This also applies to i-47.)

Cl 99 SC 99.4.3 P 43 L 20 # i-81
RAN, ADEE Intel Corporation

Comment Type TR Comment Status A

"Transmission of a verify packet is repeated if no response is received."

This "repeat" is vague, and suggests an infinite loop. It is too partial to be helpful here.

Figure 99-8 is normative and is sufficiently detailed.

SuggestedRemedy

Delete "Transmission of a verify packet is repeated if no response is received."

Response Response Status C

ACCEPT.

Cl 99 SC 99.4.3 P 43 L 27 # i-79
RAN, ADEE Intel Corporation

Comment Type E Comment Status A

"networks" appear before and after the comma. Are these different networks ("engineered closed" vs. "where it is ensured")?

Assuming the second part of the sentence explains the first, the second "networks" is unnecessary.

SuggestedRemedy

Delete "networks" after the comma.

Response Response Status C

ACCEPT IN PRINCIPLE.

"Verification may be disabled. Verification disable is intended for engineered closed networks (i.e., where it is ensured by design that the links can support preemption capability)."

Cl 99 SC 99.4.3 P 43 L 27 # i-46
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Unnecessary details: "Verification may be disabled. Verification disable is intended for engineered closed networks, networks where it is ensured by design that the components are known, in order to meet constraints on initialization time. An in-vehicle network is an example of an engineered network with constraints on initialization time."

SuggestedRemedy

Change to: The verification function may be disabled, when used in engineered closed networks.

Response Response Status C

ACCEPT IN PRINCIPLE.

The term engineered networks or closed networks is not currently defined or used in IEEE 802.3. ("Engineered links" does occur.) Therefore we should expand a little on what an engineered closed network is.

"Verification may be disabled. Verification disable is intended for engineered closed networks (i.e., where it is ensured by design that the links can support preemption capability)."

Cl 99 SC 99.4.3 P 43 L 27 # i-80
RAN, ADEE Intel Corporation

Comment Type T Comment Status A

Is the case where verification is disabled on one side but not on the other valid? The state diagrams seem to allow it, but it's not stated in the text.

SuggestedRemedy

Append to this paragraph: "Responses to verification requests from the link partner are not affected by verification disable."

Response Response Status C

ACCEPT IN PRINCIPLE. Yes, it is intended that verification works even if the link partner has disabled verification.

The term verification request is not defined.

"Verification disable does not affect the transmission of respond mPackets."

Cl 99 SC 99.4.3 P 43 L 31 # i-47
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Unnecessary separated requirements: If verification is enabled, it shall be performed as specified in Figure 99-8(a) Response shall be performed as in Figure 99-8(b).

SuggestedRemedy

Change to: The verification function shall execute the verify and response state diagrams per Figure 99-8.
Update PICS on Figure 99-8, change "Respond" to "Response" for consistency with 99.4.3

Response Response Status C

ACCEPT IN PRINCIPLE. Replace Response with Respond rather than replacing Respond with Response.

From i-45 repsonse: Search for cases of is "Response" and replace with "Respond" when referring to the packet and state machine because it is more parallel to "Verify".

Cl 99 SC 99.4.4 P 43 L 37 # i-49
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

What is this magic it: it seems to imply that it is "transmit processing" but I believe MAC Merge sublayer is meant

SuggestedRemedy

Change "it" to "MAC Merge sublayer"

Response Response Status C

ACCEPT IN PRINCIPLE. Transmit processing does the preemption.

Replace "It" with "Transmit processing"

CI 99 SC 99.4.4 P 43 L 37 # i-48
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A
 Clarify that "Transmit processing" is actually the name of the function

SuggestedRemedy

Change "Transmit processing receives" to "The Transmit Processing function (see Figure 99-3) receives"
 There are other instances in this subclause where "transmit processing" is used alone.
 Change to "Transmit Processing function" for consistency
 Similarly, in 99.4.5, change all instances of "Receive processing" and "receive processing" with "Receive Processing function"

Response Response Status C
 ACCEPT IN PRINCIPLE. Transmit processing is the name of thing. "Function" is not needed.

Transmit processing (see Figure 99-3) receives

CI 99 SC 99.4.4 P 43 L 48 # i-50
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A
 This is regurgitation of concepts already covered in this subclause: "A device can indicate that its receiver requires an additional multiple of 64 octets before preemption occurs, using the addFragSize field in the Additional Ethernet Capabilities TLV. If an additional multiple of 64 octets, addFragSize, is requested in the received Additional Ethernet Capabilities TLV, preemption does not occur until at least octets have been sent."

SuggestedRemedy

Strike text: "A device can indicate that its receiver requires an additional multiple of 64 octets before preemption occurs, using the addFragSize field in the Additional Ethernet Capabilities TLV. If an additional multiple of 64 octets, addFragSize, is requested in the received Additional Ethernet Capabilities TLV, preemption does not occur until at least octets have been sent."

Response Response Status C
 ACCEPT.

CI 99 SC 99.4.4 P 43 L 48 # i-82
 RAN, ADEE Intel Corporation

Comment Type TR Comment Status A
 The response to comment #230 against D2.2 seems to have been implemented differently than reported response. The paragraph currently in lines 43-46 was added instead of changing the following paragraph.

The result is a confusing duplication: the paragraph in lines 48-51 discusses how a device (in fact, the link partner) can set communicate its request and set addFragSize field and then how the MM sublayer should behave; but that behavior was already defined in the previous paragraph, independently of the value of addFragSize.

The specifications are usually stated for a device, not for its link partner. Stating what the link partner can do is informative in nature, so should only be a NOTE.

SuggestedRemedy

Replace the paragraph in lines 48-51 with the following NOTE:

NOTE--A device can indicate that its receiver requires an additional multiple of 64 octets before preemption occurs, using the addFragSize field in the Additional Ethernet Capabilities TLV that is sent to the link partner.

Response Response Status C
 ACCEPT IN PRINCIPLE. Delete the paragraph in lines 48-51 as requested by comment i-50 instead.

CI 99 SC 99.4.4 P 44 L 1 # i-51
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R
 "When preemption capability is active," - likely, you mean the preemption function itself

SuggestedRemedy

Change to: "When the preemption function is enabled,"

Response Response Status C
 REJECT. There is no preemption function.

The functional block that preempts is transmit processing. The preemption capability of transmit processing can be enabled or disabled.

CI 99 SC 99.4.4 P 44 L 2 # i-83
 RAN, ADEE Intel Corporation

Comment Type E Comment Status A
 Why is this text parenthesized? it seems informative, so perhaps it should be a NOTE instead?

SuggestedRemedy
 Remove the parentheses. Consider moving the enclosed sentence to a NOTE or deleting it altogether.

Response Response Status C
 ACCEPT IN PRINCIPLE. Remove the parenthesis.

CI 99 SC 99.4.4 P 44 L 2 # i-52
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
 "()" around the sentence is not needed

SuggestedRemedy
 Remove "()" around the last sentence in first para

Response Response Status C
 ACCEPT.

CI 99 SC 99.4.4 P 44 L 9 # i-53
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A
 "Transmit processing starts transmission of the remainder" - likely, "resumes" since it was interrupted before

SuggestedRemedy
 Change to "The Transmit Processing function resumes transmission of the remainder" ...

Response Response Status C
 ACCEPT IN PRINCIPLE. "Transmit processing resumes transmission of the remainder"

CI 99 SC 99.4.4 P 44 L 15 # i-112
 Scruton, Peter

Comment Type T Comment Status A PLS_CARRIER
 According to IEEE 802.3-2012 Subclause 22.2.4.1.8:

"The behavior of the CRS signal is unspecified when the duplex mode bit 0.8 in the control register is set to a logic one, as described in 22.2.4.1.8, or when the Auto-Negotiation process selects a full duplex mode of operation." This runs counter to the assertion on p44 L15 (99.4.4). Which is: "In full duplex operation, the PLS_CARRIER.indication primitive is not produced unless EEE (Clause 78) or Link Interruption (46.3.4) is supported." As a result there may be PHYs that do cause the CRS signal assertion on reception in Full Duplex.

SuggestedRemedy
 Add text to disallow the use of PHYs that will assert this signal for reasons other than the transmit media is unavailable (EEE or other). Here in 99.4.4 suggest wording change to: "The use of preemption is only allowed in full duplex operation, and the PLS_CARRIER.indication primitive shall not be produced while preemption capability is enabled by a PHY conforming to this clause, unless EEE (Clause 78) or Link Interruption (46.3.4) is supported."

Response Response Status C
 ACCEPT IN PRINCIPLE.
 MAC Merge doesn't depend on PHYs producing PLS_CARRIER.indications or not. It merely forwards the primitive to the MACs.

The editor has to check whether "In full duplex operation, the PLS_CARRIER.indication primitive is not produced unless EEE (Clause 78) or Link Interruption (46.3.4) is supported." is accurate and delete if it only applies to some PHYs.

CI 99 SC 99.4.5 P 44 L 26 # i-54
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A
 "If an mPacket containing an SMD-S is received when Receive processing was processing an incomplete preempted packet," - this statement is not clear, given that specific condition for "processing an incomplete preempted packet" is not defined here

SuggestedRemedy
 Generalize to: "If an mPacket containing an SMD-S is received while the Receive Processing function has not completed receiving the previous preempted packet," - details are included in associated SD

Response Response Status C
 ACCEPT.

Cl 99 SC 99.4.5 P 44 L 27 # i-113
Scruton, Peter

Comment Type T Comment Status A

Ambiguous: "If an mPacket containing an SMD-S is received when Receive processing was processing an incomplete preempted packet, Receive processing shall ensure that the MAC detects a FrameCheckError in that frame." Which frame does "that" refer to. In the state diagram this would refer to the previous partial.

SuggestedRemedy

Change to: "... MAC detects a FrameCheckError in the partially received frame."

Response Response Status C

ACCEPT.

Cl 99 SC 99.4.5 P 44 L 33 # i-114
Scruton, Peter

Comment Type T Comment Status A

"Other techniques may be employed to respond to a received Error control character provided that the result is that the MAC sublayer behaves as though a FrameCheckError occurred in the received frame." If this is referring to a PCS Coding error this layer should never see an Error control character. At least for 100BASE-TX for an error during frame reception the PCS should see a Code Group Error and flag RX_ER while RX_DV is still asserted and the RS underneath this layer should enforce this by handing something up that would ensure that the MAC would behave as though a FrameCheckError occurred.

SuggestedRemedy

Option A: Strike sentence.

Option B: As it is talking about enforcing a sequencing order error, we could update sentence: "Other techniques may be employed to respond to this error provided that the result is that the MAC sublayer behaves as though a FrameCheckError occurred in the received frame."

Response Response Status C

ACCEPT IN PRINCIPLE. "Error control character" should have been replaced by "an incomplete packet".
Other techniques may be employed to respond to an incomplete packet provided that the result is that the MAC sublayer behaves as though a FrameCheckError occurred in the received frame."

Cl 99 SC 99.4.5 P 44 L 36 # i-105
Law, David Hewlett Packard Enter

Comment Type T Comment Status A

I'm not sure that the description in this subclause in respect to PLS_DATA_VALID.indication being sent to the pMAC matches Figure 99-6 'Receive Processing State Diagram. The text states that 'Reception of the start of the preemptable packet begins with sending PLS_DATA_VALID.indication(DATA_VALID) to the pMAC'. According to Figure 99-6 however, pRX_DV(TRUE) is sent when rRxDV becomes true, which causes a transition from the IDLE_RX_PROC state to the pMAC_DATA_VALID state, regardless of the received packet being a express packet, a preemptable packet, or any other type of packet. The SMD_DECODE function then examines the incoming data on rRX_DATA byte by byte. If it is a preamble byte, a preamble byte is sent to the pMAC. Once a SMD is detected, anything other than a SMS-S will cause a pRX_DV(FALSE) to be sent.

Hence all packets result in preamble being sent to the pMAC, in the case of a preemptable packet this will be followed by a SFD and data. So while correct that reception of the start of the preemptable packet begins with sending PLS_DATA_VALID.indication(DATA_VALID) to the pMAC, the reception of any packet results in sending PLS_DATA_VALID.indication(DATA_VALID) to the pMAC. Since there is no description of the operation of PLS_DATA_VALID.indication sent to the eMAC in the Express filter function (see 99.4.6), and since subclause 99.4.7.1 'State diagram conventions' states that 'Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.' suggest this text simply be deleted.

SuggestedRemedy

Based on this delete the text 'Reception of the start of the preemptable packet begins with sending PLS_DATA_VALID.indication(DATA_VALID) to the pMAC.'

Response Response Status C

ACCEPT IN PRINCIPLE. Delete the sentence. The start of the next sentence could then use some context on when receive processing does it.

Replace: "Receive processing records the frame count indicated by the SMD..." with "When an SMD-S is detected, receive processing records the frame count indicated by the SMD..."

CI 99 SC 99.4.5 P 44 L 38 # i-68
 Gunther, Craig HARMAN INTERNATI

Comment Type T Comment Status A

Text here says "next fragment count" is set to 1 upon receipt of an SMD-S. 99.3.4 (pg 41, line 42) says frag_count=0 for first SMD-C. These are in conflict since TX side will set frag_count=0 in first SMD-C, yet RX side expects frag_count="next fragment count"=1 (see pg 44, line 54) for first fragment.

Figure 99-6 line 9 also shows nxtRxFrag<=0, which matches 99.3.4.

SuggestedRemedy

"next fragment count" should be set to zero, not one.

Response Response Status C

ACCEPT IN PRINCIPLE. There is also another inaccuracy here. The state machine sets nxtRxFrag to 0 between packets rather than after the SMD-S.

Delete "and sets the next fragment count to 1 for use in error checking of any subsequent mPackets for the packet." That level of detail isn't needed.

CI 99 SC 99.4.5 P 44 L 40 # i-55
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

Given the checking that is done in MAC Merge, there is no "forwarding of primitives from RS to pMAC" - it is MAC Merge that generates them

SuggestedRemedy

Change: "forwarding subsequent PLS_DATA.indication primitives from the RS to the pMAC" to "sending PLS_DATA.indication primitives to the pMAC"

Response Response Status C

ACCEPT.

CI 99 SC 99.4.5 P 44 L 43 # i-56
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status R

"that indicates the end of the packet" - it is not clear what "packet" is meant

SuggestedRemedy

Clarify whether it is the mPacket or MAC Packet

Response Response Status C

REJECT. Packet in 802.3 is a MAC packet. It is a defined term (see 1.4.312)

"packet" is correct here.

CI 99 SC 99.4.5 P 44 L 44 # i-57
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

"If they match, that indicates that the packet was preempted" - do you mean "last four octets are sent to the pMAC followed by PLS_DATA_VALID.indication(DATA_NOT_VALID)"?

SuggestedRemedy

Change "If they match" to "If the last four octets of the mPacket match the mCRC"

Response Response Status C

ACCEPT IN PRINCIPLE. For consistency, replace "they" on line 42 and line 44 with "the last four octets of the mPacket"

CI 99 SC 99.4.5 P 45 L 1 # i-58
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A

"If any of the checks do not pass," would read better as "If any of these conditions fail," Similar change for "If all the checks pass," in line 5

SuggestedRemedy

Per comment

Response Response Status C

ACCEPT IN PRINCIPLE. Change "If any of the checks do not pass," to "If any of the checks fail,"

CI 99 SC 99.4.6 P 45 L 18 # i-106
 Law, David Hewlett Packard Enter

Comment Type T Comment Status A

I'm not sure why this subclause states that the Express filter function passes the PLS_DATA.request primitive to the eMAC when it detects a mPacket containing an SMD-E. According to Figure 99-3 MAC 'Merge sublayer Functional Block Diagram' the PLS_DATA.request primitive is not connected to the Express filter function and the PLS_DATA.request primitive is related to the transmit path (see IEEE Std 802.3-2015 subclause 6.3.1.1). In addition, according to Figure 99-3, it is 'RS:PLS_DATA.request' and not 'RS PLS_DATA.request'.

SuggestedRemedy

Suggest that the text '... passes the RS PLS_DATA.indication and PLS_DATA.request to the ...' be changed to read '... passes the RS:PLS_DATA.request to the ...'.

Response Response Status C

ACCEPT.

CI 99 SC 99.4.7 P 52 L 24 # i-59
Hajduczenia, Marek Bright House Network

Comment Type ER Comment Status A
Transition crossing: RCV_R > R_MCRC_OK crossing transition from RCV_V state

SuggestedRemedy
Please do not cross state transitions

Response Response Status C
ACCEPT IN PRINCIPLE. Route the !rRxDv exit from RCV_V down between V_MCRC_OK state and R_MCRC_OK state to join the transitions going from those states to INT_EXPRESS_FILTER

CI 99 SC 99.4.7.2 P 45 L 50 # i-115
Scruton, Peter

Comment Type T Comment Status A
I understand what is intended, but I'm not a fan of the current way it is written as this is saying binary and then gives a hex.
"The binary value 0x55"
"The binary value 0xD5"

SuggestedRemedy
Change these to be 8-bit vector data <7:0> or <0:7> values. I can't seem to tell if this should be <0:7> or <7:0> it looks like rTX_DATA and pRX_DATA flips it so I think it would go in as <0:7> and then it will get flipped? I know the way it should be transmitted going down the stack it would go 1 then 0101010 and SFD as 10101011...
the 8-bit vector <0:7> of 0x55
the 8-bit vector <0:7> of 0xD5

Response Response Status C
ACCEPT IN PRINCIPLE. Ordering for the vectors is covered in the first two paragraphs of 99.4.7.4. Bit 0 of the vector maps to the first primitive.

Add after the 2nd paragraph of 99.4.7.4, "When the value of a vector is expressed as a hexadecimal number, the LSB corresponds to bit 0 and the MSB corresponds to bit 7."

For PREAMBLE and SFD constants, delete "The binary value"

CI 99 SC 99.4.7.3 P 34 L 37 # i-86
RAN, ADEE Intel Corporation

Comment Type ER Comment Status A
Variable name referred to as verifyTime in the description and in other places

SuggestedRemedy
Rename to verifyTime.

Response Response Status C
ACCEPT IN PRINCIPLE. The location is page 47 line 37. Change verify_Time to verifyTime.

CI 99 SC 99.4.7.3 P 46 L 44 # i-84
RAN, ADEE Intel Corporation

Comment Type ER Comment Status A
There is an unusual character used here for multiplication - it looks like an asterisk but is not (not found by search) and is not the cross sign which is required per style manual.

The same character is used in state diagrams where normally a plain asterisk is used.

This character is not used in other clauses, and formatting changes may cause it to change to something unreadable.

SuggestedRemedy
Search and replace all 29 occurrences of this character:

- to multiplication sign (cross character) in text and equations
- to an asterisk in state diagrams.

Response Response Status C
ACCEPT IN PRINCIPLE.

If the character is used for multiplication, change it to multiplication sign.

On page 46 line 44 the operation is a Boolean AND, not multiplication.

The table of symbols for state machines in Clause 22 has an asterisk in the Symbol font for Boolean AND. That is what the character is.

Cl 99 SC 99.4.7.3 P 46 L 52 # i-85
RAN, ADEE Intel Corporation

Comment Type TR Comment Status A

This is a long, inconsistent and confusing formula.

The + sign is used both for logical-or and for numerical addition, comparing Booleans to TRUE is inconsistent, spacing is inconsistent, and the parentheses do not prevent ambiguity (no parentheses around the numerical part).

Suggested changes aimed to remove ambiguity and improve readability.

SuggestedRemedy

Add a new variable definition:

preemptableFragSize: Boolean variable that is true when a sufficient size of the current preemptable packet has been transmitted so that it may be preempted. Value is TRUE if fragSize >= minFrag * (1+addFragSize) - 4).

In the definition of preempt, change:

"The value of preempt is: pAllow * (eTx=TRUE + hold=TRUE) * fragSize>=(minFrag * (1 + addFragSize) - 4) * MIN_REMAIN"

to

"The value of preempt is: pAllow * (eTx+hold) * preemptableFragSize * MIN_REMAIN".

Response Response Status C

ACCEPT.

Cl 99 SC 99.4.7.7 P 50 L 1 # i-116
Scruton, Peter

Comment Type T Comment Status R

Figure 99-5 Transmit Processing State Diagram

As PLS_Carrier.indication could be produced in EEE or Link Interruption perhaps it may be advisable to have an additional entrance condition in the START_PREAMBLE, and the transition from RESUME_WAIT into RESUME_PREAMBLE to also be And-ed with PLS_Carrier.indication=CARRIER_OFF. This way if a Preemptable packet arrives while the media is unavailable the decision as to whether to send this frame will not be made until after the media is available. This way if the media is unavailable the (an Express Frame may be available at that time). Related to this if EEE is allowed (looks to currently be the case) then LP_IDLE.request shall not be set to ASSERT when frames need to be transmitted and also 802.3-2012 subclause 22.7.2:

"The operation of LPI in the PHY requires that the MAC does not send valid data for a time after LPI has been de-asserted as governed by resolved Transmit Tw_sys defined in 78.4.2.3. This wake up time is enforced by the transmit LPI state diagram and the rules mapping CARRIER_SENSE.indication defined in 22.2.1.3. The implementation shall conform to the behavior described by the transmit LPI state diagram shown in Figure 22-23."

SuggestedRemedy

Solution A: Specifically allow EEE: Add signal LP_IDLE.request into Figures 99-2 and 99-3. Add necessary states and transitions to Figure 99-5 to accomplish: - Allow asserting LP_IDLE.request, but when traffic is to be sent deassert and timeout before transmit. Solution B: A statement requiring that if EEE is enabled ensure that LP_IDLE.request remains Deasserted.

Response Response Status C

REJECT. This is not needed. MAC MERGE passes any PLS_Carrier.indication received from the RS to the pMAC and the eMAC. These indications shouldn't occur while a packet is being sent. The MACs handle them.

Cl 99 SC 99.4.7.7 P 51 L 8 # i-107
Law, David Hewlett Packard Enter

Comment Type T Comment Status A

The variable 'resumeRx' is set FALSE in the state IDLE_RX_PROC of Figure 99-6 'Receive Processing State Diagram' however is never set TRUE and is never used.

SuggestedRemedy

Suggest that the assignment 'resumeRx <= FALSE' in state IDLE_RX_PROC be deleted, and that definition of the variable resumeRx in subclause 99.4.7.3 'Variables' on page 47, line 14 be deleted.

Response Response Status C

ACCEPT.

CI 99 SC 99.4.7.7 P 52 L 1 # i-117
Scruton, Peter

Comment Type T Comment Status A

Figure 99-7
States: RCV_V and RCV_R
These both have pRX_DV(False) calls. This looks to be done with the intention that if there is a V or an R saying that any continuation of a preempted frame would be wrong. I don't think the R would imply that, as 99.4 would seem to indicate that it should always be ready to accept.
It is strange that these have the affect of altering the states typically used in figure 99-6. If the intention is to discard in this case it could be done with an additional state transition in Figure 99-6.

SuggestedRemedy

Remove those pRX_DV(False) calls in Figure 99-7 states RCV_V and RCV_R.
If it is desired to discard when the remote side does V add a transition in Figure 99-6 from CHECK_FOR_RESUME to ASSEMBLY_ERROR on condition V (because entering ASSEMBLY_ERROR increments a statistic see 30.14.1.8 counter for Assembly errors it may or may not be desirable to count this as an assembly error, if not then this may be a new state with the DISCARD function inside and then a transition on lRxDV to IDLE_RX_PROC).

Response Response Status C

ACCEPT IN PRINCIPLE. These should have been changed to eRX_DV when the reception of Verify and Respond packets was moved to the Express state diagram.

In States RCV_V and RCV_R, change pRX_DV to eRX_DV.

CI 99 SC 99.4.8 P 52 L 47 # i-96
Remein, Duane Futurewei Technologie

Comment Type TR Comment Status A

Requirement without corresponding PICS: "HRT shall be no more than (1240 + 512 x addFragSize) bit times."

SuggestedRemedy

Change to factual statement "HRT is no more than (1240 + 512 x addFragSize) bit times."

Response Response Status C

ACCEPT IN PRINCIPLE. HRT defines the maximum response time so it should say: HRT is (1240 + 512 x addFragSize) bit times.

CI 99 SC 99.4.8 P 53 L 1 # i-87
RAN, ADEE Intel Corporation

Comment Type E Comment Status A

Figure 99-8 is really two figures in one frame. That's acceptable, but the formatting should clarify that these are two separate diagrams.

There seems to be enough space on both sides to make that separation.

SuggestedRemedy

Please edit to separate the two diagrams into distinct "left" and "right" sides, with some vertical white space separating them. In addition, align the labels of these two parts horizontally at the bottom.

Response Response Status C

ACCEPT.

CI 99 SC 99.4.8 P 53 L 27 # i-67
Gunther, Craig HARMAN INTERNATI

Comment Type E Comment Status A

This is purely a comment about readability. The transition from WAIT_FOR_RESPONSE to VERIFIED is based on rcv_r. The placement of the rcv_r transition would be clearer if it was moved to the right side of the transition arc rather than where it is shown now. This would put it near the bottom of the WAIT_FOR_RESPONSE box rather than above the VERIFIED box.

SuggestedRemedy

Move rcv_r transition label to the right side of the transition.

Response Response Status C

ACCEPT.

CI 99 SC 99.5.2.2 P 54 L 34 # i-60
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A

Change "IEEE Std 802.3br-2016, Clause 99" to "IEEE Std 802.3br-201X, Clause 99" - this is not a published amendment yes

SuggestedRemedy

Per comment, two instances

Response Response Status C

ACCEPT.

Cl 99 SC 99.5.3.1 P 55 L 13 # i-88
RAN, ADEE Intel Corporation

Comment Type TR Comment Status A
In item MM4, feature is "Additional Capabilities TLV". The subclause referred has no requirement from the TLV, and this is really not a function of this clause.

SuggestedRemedy
Delete MM4, possibly merge the comment into MM3.

Response Response Status C
ACCEPT IN PRINCIPLE.
Change MM3 to "Enabled only if the link partner announces support for preemption capability as described" and delete MM4

Cl 99 SC 99.5.3.1 P 55 L 16 # i-95
Remein, Duane Futurewei Technologie

Comment Type TR Comment Status A
Ambiguous language mandatory requirement in Value/Comment "Should be disabled on link failure"

SuggestedRemedy
Change to: "disabled on link failure"

Response Response Status C
ACCEPT IN PRINCIPLE. Change to "disabled on link failure indication" to reflect the text in 99.4.2

Cl 99 SC 99.5.3.1 P 55 L 16 # i-94
Remein, Duane Futurewei Technologie

Comment Type TR Comment Status A
Ambiguous language for mandatory requirement in Value/Comment "Should be sent in LLDP frame addressed to Nearest Bridge group address"

SuggestedRemedy
Change to: "Sent in LLDP frame addressed to Nearest Bridge group address"

Response Response Status C
ACCEPT IN PRINCIPLE. Remove MM4 instead. There isn't a corresponding shall anymore

Cl 99 SC 99.5.3.1 P 55 L 39 # i-61
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
No need to "shall" in PICS

SuggestedRemedy
Change "MAC Control sublayer shall not generate PAUSE" to "MAC Control sublayer does not generate PAUSE"

Response Response Status C
ACCEPT.