

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 00 SC 0 P L # 145
Regev, Alon Ixia

Comment Type E Comment Status D
Header contains the text "IEEE 802.3br Task Force name Task Force"

SuggestedRemedy
Replace "IEEE 802.3br Task Force name Task Force" with "IEEE P802.3br Interspersing Express Traffic Task Force"

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 00 SC 0 P L # 146
Regev, Alon Ixia

Comment Type E Comment Status D
Header contains the texts "IEEE Draft P802.3br/D0.1"

SuggestedRemedy
Replace "D0.1" with "D1.0" (or the current version of the draft).

Proposed Response Response Status W
PROPOSED ACCEPT IN PRINCIPLE.
Applies to Clause 1 and Clause 99.

Cl 00 SC 0 P1 L2 # 17
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D
Draft from the future !

SuggestedRemedy
Publication month is off :) It was November, not December

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 00 SC 0 P13 L0 # 69
Tretter, Albert Siemens AG

Comment Type E Comment Status D
In the header of this page the draft number and the date is not correct
"IEEE Draft P802.3br/D0.1; 4th April 2014"

This is also the case at the following pages: 14, 29 - 42, 44 - 48

SuggestedRemedy
Please correct

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 00 SC 0 P43 L # 136
Vahid Tabatabaee Broadcom

Comment Type T Comment Status A
Need to add checks for the size of non-final fragment
non-final frag size is >= 64B
non-final frag size is multiple of 8 octets

SuggestedRemedy

Response Response Status C
ACCEPT IN PRINCIPLE. From the discussion, there was consensus that the multiple of 8 octets requires extra complexity in the transmitter that can be as great or greater than the burden on the receiver for handling receiving a non-multiple of 8 data octets in a fragment. The burden on the receiver is not that large. Therefore, it was agreed that the multiple of 8 octet requirement from the baseline should be removed.

We could not come up with a reason for the receiver to enforce the minimum fragment size requirement. Make no change to fragment checks.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 00 SC 0 P 43 L 46 # 135
 Vahid Tabatabaee Broadcom

Comment Type T Comment Status A
 Need to check if cFrameCnt = rxFrameCnt and if it does not discard the frame

SuggestedRemedy

Response Response Status C
 ACCEPT IN PRINCIPLE.
 Add * cFrameCnt = rxFrameCnt to the transition from CHECK FOR RESUME to CHECK FRAG CNT
 Add + (C * cFrameCnt /= rxFrameCnt) to the transition from CHECK FOR RESUME to ASSEMBLY ERROR

Cl 00 SC 0 P 43 L 46 # 134
 Vahid Tabatabaee Broadcom

Comment Type TR Comment Status D
 Need to check if cFrameCnt matches the rxFrameCnt and if it does not match discard the frame

SuggestedRemedy
 This check should be added in Receive Processing State Diagram (same as CHECK FRAG CNT) there should be CHECK FRAME CNT

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE. See comment #135

Cl 00 SC 30.14.1.7 P 23 L 28 # 87
 Tretter, Albert Siemens AG

Comment Type E Comment Status D
 aMACMergeFrameSmdErrorCount:
 ...unknown SMD value or arriving with an SMD_C when no frame is in progress.

SuggestedRemedy
 Typo: SMD_C should be changed to SMD-C as it is named in the rest of the document.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 01 SC 1.3 P 14 L 10 # 147
 Regev, Alon Ixia

Comment Type E Comment Status D
 In the editor's note, in the sentence "Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks s the current title of IEEE 802.1Q", the verb "is" is written as "s"

SuggestedRemedy
 correct "s" to "is"

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 01 SC 1.4 P 14 L 18 # 19
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D
 Numbers for individual definitions should be provided:
 express Media Access Control (eMAC)
 preemptable Media Access Control
 MAC Merge sublayer

SuggestedRemedy
 Per comment. Also remove empty definition in line 24.

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 01 SC 1.4 P 14 L 18 # 149
 Regev, Alon Ixia

Comment Type T Comment Status D
 In the definition of express Media Access Control and preemptable Media Access Control, we are using the text "which is the client of a MAC Merge sublayer that handles XXX frames" (where XXX is express or preemptable). As there is only 1 MAC merge sublayer, both the eMAC & pMAC are clients of the same MAC Merge sublayer, so the pMAC & eMAC cannot be distinguished by the MAC merge sublayer of which they are clients.

SuggestedRemedy
 Change "which is the client of a MAC Merge sublayer that handles" to "which is the client of a MAC Merge sublayer service interface that handles"

Proposed Response Response Status W
 PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 01 SC 1.4 P 14 L 21 # 37
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
preemptable Media Access Control could use an acronym

SuggestedRemedy
change "preemptable Media Access Control" to "preemptable Media Access Control (pMAC)"

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 01 SC 1.4 P 14 L 21 # 151
Regev, Alon Ixia

Comment Type E Comment Status D
The inclusion of the abbreviation is inconsistent between the definition of "express Media Access Control" and "preemptable Media Access Control" (express Media Access Control includes "(eMAC)" in the definition in section 1.4).

SuggestedRemedy
change "preemptable Media Access Control" to "preemptable Media Access Control (pMAC)" in the definition.

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 01 SC 1.4 P 14 L 24 # 124
Jones, Peter Cisco

Comment Type E Comment Status D
Remove blank placeholder definition "1.4.x" between "preemptable Media Access Control" and "MAC Merge sublayer"

SuggestedRemedy
remove text

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 01 SC 1.4 P 14 L 26 # 150
Regev, Alon Ixia

Comment Type E Comment Status D
It would be useful to have the definition of MAC Merge sublayer refer to the clause where it is described

SuggestedRemedy
Add reference "(See IEEE Std 802.3, Clause 99)" to the MAC Merge Sublayer definition in section 1.4

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 01 SC 1.4 P 14 L 27 # 152
Regev, Alon Ixia

Comment Type E Comment Status D
The definition of "MAC Merge sublayer" is referencing "express traffic" and "preemptable traffic", but these terms are not defined anywhere

SuggestedRemedy
Add the following definitions to section 1.4:

1.4.x Express Traffic: Traffic frames transmitted through an express Media Access Controller.

1.4.x preemptable Traffic: Traffic frames transmitted through a preemptable Media Access Controller.

Proposed Response Response Status W
PROPOSED ACCEPT IN PRINCIPLE.
Controller should be Control

Cl 01 SC 1.4.x P 14 L 21 # 137
Mitsuru, Iwaoka Yokogawa Electric Cor

Comment Type E Comment Status D late
It is better to add an abbreviation "pMAC" as same as the eMAC.

SuggestedRemedy
Insert "(pMAC)" after "preemptable Media Access Control".

Proposed Response Response Status W
PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 01 SC 1.5 P 14 L 36 # 148
 Regev, Alon Ixia
 Comment Type E Comment Status D
 In the text "preemptable Medica Access Control", "Medica" should be "Media"
 SuggestedRemedy
 change "Medica" to "Media"
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 30 SC P 15 L 3 # 153
 Regev, Alon Ixia
 Comment Type E Comment Status D
 Extra "[" appears
 SuggestedRemedy
 Remove the "["
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 01 SC 1.5 P 14 L 37 # 143
 Mitsuru, Iwaoka Yokogawa Electric Cor
 Comment Type E Comment Status D late
 It is better to specify following abbreviations:
 - HRT
 - MMSI
 SuggestedRemedy
 Insert following abbreviations:
 HRT hold response time
 MMSI MAC Merge Service Interface
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 30 SC 30.12.2.1.30 P 20 L 20 # 73
 Tretter, Albert Siemens AG
 Comment Type E Comment Status D
 clause "..... supports preemption.capability"
 It is not clear if the word "capability" is really necessary.
 SuggestedRemedy
 Proposal:
 A read-only Boolean value used to indicate whether the given port (associated with the local system) supports preemption;
 OR
 A read-only Boolean value used to indicate whether the given port (associated with the local system) supports preemption capability;
 Proposed Response Response Status W
 PROPOSED REJECT. It is needed. See comment 73

Cl 1.4 SC 1.4x P 14 L 24 # 70
 Tretter, Albert Siemens AG
 Comment Type E Comment Status D
 Subclause "1.4x" is without any text
 SuggestedRemedy
 Please delete
 Proposed Response Response Status W
 PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 30 SC 30.12.2.1.33 P 20 L 47 # 129
 Vahid Tabatabaee Broadcom

Comment Type T Comment Status D

aLldpXdot3LocAddFragSize is defined to specify the size of non-final fragment. However, in page 26 (table 79-7a) and section 99.4.4 there is no distinction between final and non-final fragments.

SuggestedRemedy

A 3-bit integer value used to indicate, in units of 64 octets, the minimum number of octets over 64 octets required in non-final fragments by the receiver on the given port associated with the local system

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 30 SC 30.12.3.1.27 P 21 L 38 # 130
 Vahid Tabatabaee Broadcom

Comment Type T Comment Status D

aLldpXdot3RemAddFragSize is defined to specify the size of non-final fragment. However, in page 26 (table 79-7a) and section 99.4.4 there is no distinction between final and non-final fragments.

SuggestedRemedy

A 3-bit integer value used to indicate, in units of 64 octets, the minimum number of octets over 64 octets required in non-final fragments by the receiver on the given port associated with the local system

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 30 SC 30.14.1.1 P 22 L 7 # 48
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D discuss

Wording could be improved to avoid complex statements: "The value of aMACID for the express oMACEntity is assigned so as to uniquely identify an oMACMerge entity among the subordinate managed objects of the containing object.;"

SuggestedRemedy

Change to "The value of aMACID assigned to the express oMACEntity uniquely identifies oMACMerge entity among the subordinate managed objects of the containing object.;"

Proposed Response Response Status W

PROPOSED ACCEPT.
 But discuss because the current language is what other a__ID attributes use.

Also, this attribute is not needed at all - the containment identifies the express oMACEntity associated with the oMACMergeEntity. The attribute was created before the containment relationship was worked out.

Delete subclause 30.14.1.1 aMACMergeID

Cl 30 SC 30.14.1.4 P 22 L 39 # 103
 Tretter, Albert Siemens AG

Comment Type T Comment Status A Discuss

aMACMergeStatusEnable => BEHAVIOUR DEFINED AS:
 "The status of the MAC Merge function may be modified to the indicated value via a SET operation."

Is it really necessary to provide the ability to disable the MAC Merge function via a set operation.

Is it possible to stop the MAC Merge function in the middle of a preempted frame?

SuggestedRemedy

Discussion needed.

Response Response Status C

ACCEPT IN PRINCIPLE. There are other enables that allow a capability to enable and disable a function. 30.3.1.1.6 aMACEnableStatus can disable the whole MAC. Other Enable attributes don't specify what happens to a frame during the transtion so this attribute is valid as it is.

In the case of this object, it is clear from the state machines what happens because disabling changes the value of a variable and no frame loss or corruption happens due to the transition.

No change needed

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 30 SC 30.14.1.4 P 22 L 39 # 102
Tretter, Albert Siemens AG

Comment Type T Comment Status R
aMACMergeStatusEnable => APPROPRIATE SYNTAX
value "unknown" => transmit preemption status is unknown

How can a system be in this state?

SuggestedRemedy
It should be named "not supported"

Response Response Status C
REJECT.
Other similar status attributes allow for the possibility that management is unable to determine the status.
See aDuplexStatus which can take the values
half duplex Half duplex mode
full duplex Full duplex mode
unknown Duplex status unknown

CI 30 SC 30.14.1.4 P 22 L 40 # 126
Jones, Peter Cisco

Comment Type T Comment Status D
Why do we need the value of unknown for aMACMergeStatusEnable. It's either on or off.
What effect would setting to unknown have?

SuggestedRemedy
Remove the unknown value.

Proposed Response Response Status W
PROPOSED REJECT. See comment # 102

CI 30 SC 30.14.1.5 P 23 L 2 # 26
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D
description of "verifying" value should be indented accordingly, so that the end of it (in line 3) does not look like a new value on its own.

SuggestedRemedy
Indent the text "before being activated" to align it with the beginning of the word "transmit" one line above

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 30 SC 30.14.1.6 P 23 L 18 # 156
Regev, Alon Ixia

Comment Type E Comment Status D
Reference to section 99.4.7.7 should indicate that it's a reference (i.e. change "(99.4.7.7)" to "(see 99.4.7.7)")

SuggestedRemedy
change "(99.4.7.7)" to "(see 99.4.7.7)"

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 30 SC 30.14.1.6 P 23 L 18 # 49
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R
Reference to state diagram can be simplified: "ASSEMBLY ERROR state of the Receive Processing State Diagram is entered (99.4.7.7)."

SuggestedRemedy
Change to read: "ASSEMBLY ERROR state is entered (see Figure 99-5)."

Response Response Status C
REJECT.
For a reference between Clauses, it is preferable to reference the section number so the table of contents can get the reader of a paper copy to the correct page.

CI 30 SC 30.14.1.7 P 23 L 28 # 50
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A
Statement is not complete: "Incremented each time the BAD FRAG state of the Receive Processing State Diagram is entered.99.4.7.7);"

SuggestedRemedy
Change to read: "The counter is incremented each time the BAD FRAG state is entered (see Figure 99-5).;"

Response Response Status C
ACCEPT IN PRINCIPLE. Leave as a subclause reference

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 30 SC 30.14.1.7 P 23 L 29 # 157
 Regev, Alon Ixia
 Comment Type E Comment Status D
 Reference to 99.4.7.7 should include an opening parenthesis as well as "see".
 SuggestedRemedy
 Change "99.4.7.7)" to "(see 99.4.7.7)"
 Proposed Response Response Status W
 PROPOSED ACCEPT.

CI 30 SC 30.14.1.8 P 23 L 39 # 27
 Hajduczenia, Marek Bright House Network
 Comment Type E Comment Status D
 "attributes indicates" - the proper word is "represents"
 SuggestedRemedy
 change "indicates" to "represents"
 Proposed Response Response Status W
 PROPOSED ACCEPT.

CI 30 SC 30.14.1.8 P 23 L 39 # 105
 Tretter, Albert Siemens AG
 Comment Type T Comment Status A
 aMACMergeFrameAssOkCount:
 The sum of aMACMergeFrameAssOkCount and aMACMergeFrameAssErrorCount attributes indicates the total number of received fragmented MAC frames.
 (1)In which state is the counter for the object "aMACMergeFrameAssOkCount" located?
 (2)Question to the statement that the sum of aMACMergeFrameAssOkCount and aMACMergeFrameAssErrorCount attributes indicates the total number of received fragmented MAC frames:
 - What about the fragments discarded due to wrong SMD (SMD-C instead of SMD-S)?
 is not completely correct. Because of the problem we have with the defragmentation problem
 SuggestedRemedy
 Clarification needed
 Response Response Status C

CI 30 SC 30.14.1.9 P 23 L 49 # 88
 Tretter, Albert Siemens AG
 Comment Type E Comment Status A Discuss
 aMACMergeFrameExpressCountTx:
 A count of MAC frames transmitted by the express MAC
 Are the frames from the eMAC not counted in the eMAC itself?
 The e MAC is not part of the MAC Merge layer.
 SuggestedRemedy
 I assume it is ment that this counter counts the express frames received from the eMAC.
 Response Response Status C
 ACCEPT IN PRINCIPLE.
 Yes, this counter is needed since each MAC has a count of how many frames it has transmitted. The counter was in the accepted baseline for management when we hadn't yet developed the containment model.
 Delete the counter

ACCEPT IN PRINCIPLE.
 The counter isn't tied to a state. Perhaps you are pointing out that the description might indicate when the counter is incremented. There is no direct state machine description for when to increment this counter as the same state is entered when a fragmented frame completes successfully and when an unfragmented frame completes.
 Usually we do the state machine description for error conditions as it is the easiest way to describe them clearly.
 No change needed for (1)
 For (2) this is the closest the receiver can come to knowing how many frames were fragmented but as the commenter points out, it is approximate since some errored fragmented frames (e.g. not receiving the start of the initial fragment) won't be counted. Since the text is informative, delete it.

CI 30 SC 30.14.1.9 P 23 L 49 # 51
 Hajduczenia, Marek Bright House Network
 Comment Type T Comment Status D
 "A count of MAC frames transmitted by the express MAC.;" - we have a nice acronym for "express MAC"
 SuggestedRemedy
 Change to read: "A count of MAC frames transmitted by eMAC.;"
 Same change on page 24, line 5
 Proposed Response Response Status W
 PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 30 SC 30.14.1.xx P 24 L 40 # 104
Tretter, Albert Siemens AG

Comment Type T Comment Status A discuss

I miss an object indicating the number of octets of the minimal non preemtable frame length which was negotiated at startup between both link partners.

SuggestedRemedy

The proposed object should be added.

Response Response Status C

ACCEPT IN PRINCIPLE.

There is no negotiation. The transmitter uses the value the link partner's receiver sent.

Since the TLV might not continue to be sent once preemption is active, add an attribute for aAddFragSize. Should there also be an attribute for the link partner addFragSize (the one requested of the link partner transmitter)?

Cl 30 SC 30.2.2.1 P 15 L 16 # 121
Jones, Peter Cisco

Comment Type T Comment Status A

The text added to the end of the definition for "oMACEntity" doesn't make sense.

Based on "Figure 30-3— DTE System entity relationship diagram" below, I think this is supposed to say
"If oMACMergeEntity is implemented, the oMACEntity for the express MAC contains a instance of oMACMergeEntity, which then contains an oMACEntity for the pMAC."

SuggestedRemedy

Review current text and text ion comment above, correct as required.

Response Response Status C

ACCEPT IN PRINCIPLE.

"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 preemtable MAC (pMAC)."

Cl 30 SC 30.2.2.1 P 15 L 17 # 38
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

Newly added text does not read right: "If oMACMergeEntity is implemented, an oMACEntity for the express MAC. IfoMACMergeEntity is implemented, an oMACEntity for the pMAC, is contained in the oMACMergeEntity."

SuggestedRemedy

Implement changes shown in >><<

If oMACMergeEntity is implemented, >>oMACEntity represents<< the express MAC. If>> <<oMACMergeEntity is implemented, an oMACEntity for the pMAC>> <<is contained >>within<< the oMACMergeEntity.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. See Comment #121

Cl 30 SC 30.2.2.1 P 15 L 17 # 138
Mitsuru, Iwaoka Yokogawa Electric Cor

Comment Type E Comment Status D late

The inserted text contains several grammatical errors.

SuggestedRemedy

Modify the inserted text as follows:

If oMACMergeEntity is implemented, an oMACEntity for the express MAC contains the oMACMergeEntity, and an oMACEntity for the pMAC is contained in the oMACMergeEntity.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. See Comment #121

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 30 SC 30.2.2.1 P 15 L 26 # 39
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

Text of the description for oMACMergeEntity could be simplified for clarity without loss of information

SuggestedRemedy

Revise as follows:

If oMACMergeEntity is implemented, oMACMergeEntity is represented by two oMACEntity objects: one for an express MAC (eMAC) and one for a preemptable MAC (pMAC) (see Clause 99). oMACMergeEntity managed object class provides the management controls necessary for the MAC Merge sublayer.

Response Response Status C

ACCEPT IN PRINCIPLE.

If implemented, a single instance of oMACMergeEntity is contained within oMACEntity object for an express MAC (eMAC) and contains an instance of oMACEntity object for a preemptable MAC (pMAC) (see Clause 99). oMACMergeEntity managed object class provides the management controls necessary for the MAC Merge sublayer.

CI 30 SC 30.2.2.1 P 15 L 26 # 122
 Jones, Peter Cisco

Comment Type T Comment Status A

The definition for "oMACMergeEntity" says "a single instance of oMACMergeEntity is contained within two oMACEntity objects: one for an express MAC, eMAC; and one for a preemptable MAC, pMAC" which contradicts to "Figure 30-3— DTE System entity relationship diagram" below. I think oMACMergeEntity contains an oMACEntity for the pMAC.

SuggestedRemedy

Review definitions and containment model and correct inconsistencies.

Response Response Status C

ACCEPT IN PRINCIPLE. See Comment # 39

CI 30 SC 30.2.2.1 P 15 L 26 # 139
 Mitsuru, Iwaoka Yokogawa Electric Cor

Comment Type T Comment Status D late

The containment relationship between oMACMergeEntity and two oMACNetity objects specified here differs from the containment relation ship specified in the Figure 30-3 (see 30.2.3).

SuggestedRemedy

Modify the first sentence as follows:

If implemented, a single instance of oMACMergeEntity is contained within oMACEntity object for an express MAC (eMAC) and contains oMACEntity object for a preemptable MAC (pMAC)(see Clause 30.2.3).

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. See Comment # 39

CI 30 SC 30.2.2.1 P 15 L 5 # 20
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Editorial instruction is not correct: "Replace the description of oMACEntity with the following:"

SuggestedRemedy

Change the editorial instruction to read: "Modify the description of oMACEntity as follows:"

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

CI 30 SC 30.2.3 P 17 L 1 # 140
 Mitsuru, Iwaoka Yokogawa Electric Cor

Comment Type E Comment Status D late

In Figure 30-3, a line between oOAM and oMACEntity shall denote one-to-one relationship.

SuggestedRemedy

Add arrow to the line between oOAM and oMACEntity.

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 30 SC 30.2.3 P 17 L 1 # 21
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Consider adding red boxes around changes in the said figure to make sure that everybody is aware of the necessary modifications to the figure.

SuggestedRemedy

Mark the said red boxes as Editorial Notes to be removed prior to publication. They will be very helpful during the review process.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. Either 2 red boxes or a red polygon line will be used to outline the added material.

Cl 30 SC 30.2.3 P 17 L 1 # 141
Mitsuru, Iwaoka Yokogawa Electric Cor

Comment Type E Comment Status D late

In Figure 30-3, oMACMergeEntity is specified in 30.14.1, not 30.3.9.

SuggestedRemedy

Replace "30.3.9" with "30..14.1".

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. 30.14.1

Cl 30 SC 30.2.3 P 17 L 48 # 154
Regev, Alon Ixia

Comment Type E Comment Status D

We should not be using the name "preemptive MAC". We should only be using preemptable MAC or express MAC

SuggestedRemedy

replace "pMAC = Preemptive MAC" with "pMAC = Preemptable MAC"

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. Delete instead (see comment # 40)

Cl 30 SC 30.2.3 P 17 L 48 # 40
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

pMAC=Preemptive MAC - this is already defined. No need to repeat

SuggestedRemedy

Remove "pMAC=Preemptive MAC"

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 30 SC 30.2.5 P 18 L 19 # 41
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

There are no changes to Table 30-8 shown at this time.

SuggestedRemedy

Remove Table 30-8

Proposed Response Response Status W

PROPOSED REJECT. Table 30-8 is a new table for the package for the objects for the new TLV - that is why the editing instruction says Insert Table 30-8. A new table was started for additional packages because Table 30-7 was too wide to add an additional column.

Cl 30 SC 30.2.5 P 18 L 40 # 155
Regev, Alon Ixia

Comment Type T Comment Status D

In table 30-8, the Capabilities aLldpXdot3LocPreemptSupported, aLldpXdot3LocPreemptEnabled, aLldpXdot3LocPreemptActive, and aLldpXdot3LocAddFragSize are repeated twice. I believe the second set should reflect the remote (rather than the local) capabilities

SuggestedRemedy

Rename the second "aLldpXdot3LocPreemptSupported" to "aLldpXdot3RemPreemptSupported"

Rename the second "aLldpXdot3LocPreemptEnabled" to "aLldpXdot3RemPreemptEnabled"

Rename the second "aLldpXdot3LocPreemptActive" to "aLldpXdot3RemPreemptActive"

Rename the second "aLldpXdot3LocAddFragSize" to "aLldpXdot3RemAddFragSize"

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 30 SC 30.2.5 P 18 L 40 # 125
 Jones, Peter Cisco

Comment Type E Comment Status D

Table 30–8—LLDP capabilities (additional packages) has copy and paste errors, with two copies of the aLldpXdot3Locxxx attributes, the second set should be aLldpXdot3Remxxx

SuggestedRemedy

Fix table

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.
 In the last 4 lines of the table, replace Loc with Rem

Cl 30 SC 30.2.5 P 18 L 5 # 71
 Tretter, Albert Siemens AG

Comment Type E Comment Status D

The table 30-9 is referenced twice.

SuggestedRemedy

One of the references should be to table 30-8?

Proposed Response Response Status W

PROPOSED ACCEPT. Yes, the first instance should be Table 30-8

Cl 30 SC 30.2.5 P 19 L 1 # 22
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Table 30–9 uses incorrect font format for individual entries:

- a M A C M e r g e I D
- aMACMergeSupport
- aMACMergeStatusVerify
- aMACMergeStatusEnable
- aMACMergeStatusTx
- a M A C M e r g e F r a m e AssErrorCount
- aMACMergeFrameSmdErrorCount
- aMACMergeFrameAssOkCount
- aMACMergeFrameExpressCountTx
- aMACMergeFrameExpressCountRx
- aMACMergeFragCountRx
- aMACMergeFragCountTx
- aMACMergeHoldCount

SuggestedRemedy

Please apply proper formatting per 802.3 template

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 30 SC 30.2.5.0.1 P 19 L 1 # 72
 Tretter, Albert Siemens AG

Comment Type E Comment Status D

Table 30-9 is somehow placed in the description of the aLldpXdot3PortConfigTLVsTxEnable object. I assume this table should be placed before or after this description.

SuggestedRemedy

Please correct

Proposed Response Response Status W

PROPOSED REJECT.
 Tables float to where they fit

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 30 SC 30.2.5.0.1 P 19 L 50 # 107
 Jones, Peter Cisco

Comment Type E Comment Status X

Editor's Note suggests "this should be changed to something like, "Starting with the first bit in the string, the mapping of bits to TLVs is: followed by a list of the TLVs]". I agree with this.

SuggestedRemedy

Text similar to the following

Starting with the first bit in the string, the mapping of bits to TLVs is:
 bit 0: MAC/PHY configuration/status TLV transmit is enabled
 bit 1: Power via MDI TLV transmit is enabled
 bit 2: deprecated Link Aggregation TLV transmit is enabled
 bit 3: Maximum Frame Size TLV transmit is enabled
 bit 4: EEE TLV transmit is enabled
 bit 5: Additional Ethernet Capabilities TLV is enabled.

Proposed Response Response Status W

Discuss with Adam Healy

Cl 30 SC 30.21.5.0.1 P 18 L 46 # 173
 NoName

Comment Type T Comment Status D Rouge comment

It turns out there were two EEE TLVs missing from the table and bits for them are being added in 802.3-Rev.

Also the subclause number is incorrect. It should be 30.12.1.1.1

SuggestedRemedy

Correct the base text to match 802.3-Rev, underscore changes from that base text, use the seventh bit for Additional Ethernet Capabilities TLV and correct the subclause number.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 79 SC 79.3 P 25 L 10 # 23
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Remove empty line 9/10

SuggestedRemedy

Please scrub the whole draft and remove all empty lines added without any need.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Cl 79 SC 79.3 P 25 L 12 # 142
 Mitsuru, Iwaoka Yokogawa Electric Cor

Comment Type E Comment Status D late

Table 1 is not exist in clause 79. It is Table 79-1.

SuggestedRemedy

As in comment.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 79 SC 79.3 P 25 L 13 # 24
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Wrong table number - it is Table 1 and should be Table 79-1

SuggestedRemedy

Per comment

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 79 SC 79.3 P 25 L 16 # 42
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

Wrong reference in the last column in Table 79-1

SuggestedRemedy

Is "Table 79.3.6" and should be "79.3.6"

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 79 SC 79.3 P 25 L 7 # 158
 Regev, Alon Ixia

Comment Type E Comment Status D

in the editing instructions, space is missing between "in" and "Table 1"

SuggestedRemedy

add space between "in" and "Table 1"

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 79 SC 79.3.6.1 P 25 L 42 # 25
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D
Wrong reference: is "Table 79-7" and should be "Table 79-7a"

SuggestedRemedy

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 79 SC 79.3.6.1 P 25 L 45 # 43
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R
"The bits reserved for future standardization shall be set to zero and ignored on receipt. Reserved octets shall not be transmitted and shall be ignored on receipt. Any octets not received shall be treated as all zero." - it is not clear what "Any octets not received shall be treated as all zero." really means in this context. The format of the TLV is defined. Is there a situation where some octets are not transmitted and then assumed to be transmitted on receive side? Taht is just bad engineering

SuggestedRemedy
Remove "Any octets not received shall be treated as all zero."

Response Response Status C
REJECT. The TLV was designed so that it can be extended in later versions to support additional capabilities. That is why it has a variable rather than fixed length. The future octets that may be added to support that are all reserved. Something not supporting capabilities using later octets doesn't need to send them (and might not even know that the TLV had been extended with later additions). This text is to support that.

CI 79 SC 79.3.6.1 P 25 L 48 # 108
Jones, Peter Cisco

Comment Type E Comment Status D
Remove the editors note regarding larger min fragment size, text is in the draft.

Draft text - [Editor's Note (to be removed prior to publication) - There has been a suggestion to allow for a receiver to request a larger minimum fragment size. If the task group agrees to that, a value for minimum fragment size, perhaps two or three bits indicating size as a multiple of 64, would be added here.

SuggestedRemedy
reomve editors note.

Proposed Response Response Status W
PROPOSED ACCEPT.

TYPE: TR/technical required ER/editorial required GR/general required T/technical E/editorial G/general
COMMENT STATUS: D/dispatched A/accepted R/rejected RESPONSE STATUS: O/open W/written C/closed Z/withdrawn
SORT ORDER: Clause, Subclause, page, line

CI 79 SC 79.3.6.2 P 25 L 48 # 159
Regev, Alon Ixia

Comment Type E Comment Status D
There is an editor's note indicating that an additional fragment size field may be needed in the TLV. As this field was added, the editor's note can be removed.

SuggestedRemedy
Remove the editor's note

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 79 SC 79.3.6.2 P 26 L 1 # 44
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
Title of Table 79-7a is missing

SuggestedRemedy
Suggest: Allocation of bits in Additional Ethernet Capabilities TLV

Proposed Response Response Status W
PROPOSED ACCEPT IN PRINCIPLE.
The title should be:
Additional Ethernet capabilities/status which is consistent with the other table names.
(see comment #109)

CI 79 SC 79.3.6.2 P 26 L 1 # 109
Jones, Peter Cisco

Comment Type E Comment Status D
Add missing title to table "Table 79-7a—", I think it should be "Additional Ethernet capabilities/status"

SuggestedRemedy
Fix table title.

Proposed Response Response Status W
PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 79 SC 79.3.6.2 P 26 L 1 # 45
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R
All other TLVs are defined from the top bit (15) downwards and this one TLV adopts an inverse convention.

SuggestedRemedy
Align to other TLVs, and start definition bits from 15 going towards 0.

Response Response Status C
REJECT.
The existing TLV tables are not consistent. Table 79-4 has the highest numbered bits on top. Tables 79-2, 79-3 and 79-7 have the 0 bit on top.

Since 3 of the 4 similar tables put bit 0 on top, this table does the same.

Cl 79 SC 79.3.6.2 P 26 L 1 # 46
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R
Some of features in Table 79-7a have confusing description

SuggestedRemedy
Change description of bit 0 to "Preemption support"
Change description of bit 1 to "Preemption status"
Change description of bit 3-5 to "Preemption fragment size (min)". Remove "additional fragment size needed by the receiver" from third column. Insert the following map and assign proper values to individual bit map values:
bit 3 4 5 fragment size

```
0 0 0 64
0 0 1
0 1 0
0 1 1
1 0 0
1 0 1
1 1 0
1 1 1
```

Response Response Status C
REJECT.
The editor found it confusing when "preemption" in some cases meant the capability of preempting and in other cases meant the act of preempting. For clarity, "preemption capability" is used when referring to the ability to preempt and "preemption" without capability is used for the act of preempting.

Additional fragment size is clear and consistent with the usage of the field in Clause 99.

Cl 79 SC 79.4.2 P 26 L 29 # 160
Regev, Alon Ixia

Comment Type T Comment Status D
Table 2 is missing the translation from TLV data to the LLDP objects.

For example, bits 3-5 indicate the "additional fragment size", but there is no indication that this TLV field sets the management object aLldpXdot3LocAddFragSizevariable or the variable "addFragSize" in the state machine or (see 30.12.2.1.33)

SuggestedRemedy
- In section 79.4.2, add the following data into Table 2:

TLV Variable	LLDP Local System Group managed object class attribute
preemption capability support	aLldpXdot3LocPreemptSupported
preemption capability status	aLldpXdot3LocPreemptEnabled
preemption capability active	aLldpXdot3LocPreemptActive
additional fragment size	aLldpXdot3LocAddFragSize

In section

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 79 SC 79.5 P 27 L 21 # 47
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A
PICS are missing

SuggestedRemedy
Insert PICS entries for the following shall statements:
a) The additional Ethernet capabilities field shall contain a bitmap that identifies the support and current status of additional Ethernet capabilities on the local IEEE 802.3 LAN station.
b) The bits reserved for future standardization shall be set to zero and ignored on receipt.
c) Reserved octets shall not be transmitted and shall be ignored on receipt.
d) An LLDPDU should contain no more than one Additional Ethernet Capabilities TLV

Response Response Status C
ACCEPT. Also produce the Clause 99 PICS

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99 P 6 L 13 # 18
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D
 "Task Force name" is not populated in many locations in the draft - 50 hits altogether

SuggestedRemedy
 Please use the proper name of the Task Force.

Proposed Response Response Status W
 PROPOSED ACCEPT.

CI 99 SC 99.1 P 29 L 15 # 28
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D
 Missing space in "frames.The MMSI"

SuggestedRemedy
 Change to "frames. The MMSI"

Proposed Response Response Status W
 PROPOSED ACCEPT.

CI 99 SC 99.1 P 29 L 15 # 29
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status A
 "transmission of preemptable frames can be held" - this language sounds just odd - "hold" is likely an incorrect word in here, and "suspend" should be used instead.

SuggestedRemedy
 Change to "transmission of preemptable frames may be suspended"
 Similarly, change "This clause also specifies a MAC Merge Service Interface (MMSI) providing a primitive that can hold or release transmission of preemptable frames." to read "This clause also specifies a MAC Merge Service Interface (MMSI) providing a primitive that suspends transmission of preemptable frames."

Response Response Status C
 ACCEPT IN PRINCIPLE.
 "This clause also specifies a MAC Merge Service Interface (MMSI) providing a primitive that suspends and resumes transmission of preemptable frames."

CI 99 SC 99.1 P 29 L 20 # 52
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R
 The logic is inverted in this sentence. It is not the presence of express traffic that causes suspension of preemptable traffic, but MMSI primitive that puts preemptable traffic on hold and then express traffic can be trasmitted.

SuggestedRemedy
 Change lines 20 - 25 to read: "When the preemption function is enabled and the MMSI primitive is asserted, transmission of preemptable traffic from pMAC is suspended as long as the MMSI primitive remains asserted. This allows eMAC to transmit any queued express traffic. Transmission of preemptable traffic is resumed once the MMIS primitive is deasserted."

Response Response Status C
 REJECT.
 The presence of express traffic does cause suspension of preemptable traffic. Preemption can be used without scheduled traffic. When there is no schedule to indicate ahead of time that express traffic is to arrive, the MMSI primitive doesn't provide a latency advantage. In that case, preemption works fine without use of the MMSI primitive.

The MMSI primitive enables lower latency when 802.1Qbv Enhancements for Scheduled Traffic is used. The MSSSI can be used to provide lower latency for scheduled traffic by asserting hold based on the schedule. This allows preemption to occur before the scheduled traffic is ready to transmit so it experiences no preemption latency.

This is all what was adopted in the baseline.

CI 99 SC 99.1 P 29 L 22 # 161
 Regev, Alon Ixia

Comment Type E Comment Status D
 In the sentence "When preemption is active, MAC Merge allows frames provided over an express MAC service interface (express traffic) to the eMAC to interrupt transmission of a preemptable frame being transmitted by the pMAC.", the "to the eMAC" should be "from the eMAC" as the Tx frames are received by the MAC merge sublayer from the eMAC

SuggestedRemedy
 change "to the eMAC" to "from the eMAC"

Proposed Response Response Status W
 PROPOSED REJECT.
 This is talking about what happens at the eMAC service interface (i.e. the interface above the MAC) - at that interface a transmitted frame is going to the eMAC

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.1 P 29 L 27 # 12
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status R

Text in lines 27-31 describes the operation with preemption disabled. In this case, it seems that MAC Merge does nothing more than add more queues to the system, allowing clients to transmit through either path, though without preemption option. I believe this is unnecessary and does not fit into the objectives of this TF.

SuggestedRemedy

Remove lines 27-31. Insert the following text: "When the preemption function is disabled, eMAC is disabled and MAC Merge exchanges data only with pMAC."

Response Response Status C

REJECT.

This behavior was strongly supported during discussions that lead to formulating the baseline. It is in the baseline. It provides a benefit.

It enables the MAC Client to see a consistent interface regardless of whether preemption is enabled or not. This simplifies the transtion to preemption active.

It isn't adding queues - there is no queue below the MAC Client interface and the number of service classes supported above the MAC Client interface is independent of whether there are two interfaces or one.

CI 99 SC 99.1 P 29 L 36 # 13
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status A

"When the MAC Merge sublayer and the Time Synchronization Service Interface are both supported, the MAC Merge sublayer acts as the generic Reconciliation sublayer (gRS) as defined in 90.5"

SuggestedRemedy

This topic was not discussed within 802.3 in a broader fashion, nor 802.3bf participants have been engaged. For this statement to work, we would need to (a) explore what the problem really is, (b) make changes to Clause 90 and TSSI definition, and not create confusion with two locations where TSSI is hooked and defined in two different places.

Suggest to:

- a) remove lines 36-42
- b) remove TSSI from Figure 99-1
- c) make necessary changes to Clause 90, showing TSSI with MAC Merge
- d) once that is done and agreed with, add subclause (not just a brief statement) to Clause 99 showing relationship with TimeSync and how to build a PHY with TimeSync and MAC Merge in place.

Response Response Status C

ACCEPT IN PRINCIPLE.

Broader 802.3 review can occur during Working Group ballot - that is what Working Group ballot is for.

The TSSI is broken without this change because there are cases where one will not be able to tell whether an indication on that interface is from the frame on the eMAC or on the pMAC.

Also, only the MAC Merge Sublayer sees SFDs on the pMAC path when preemption is active. The descriptions in Clause 90 apply equally well to the MAC Merge sublayer as to the various generic MII sublayers.

The gRS is not part of the PHY.

To make this more clear to a reader of Clause 90, add to the end of the first paragraph of 90.5.

When a MAC Merge sublayer (see Clause 99) is present, the term generic Reconciliation Sublayer (gRS) is used to denote the MAC Merge sublayer rather than the RS Sublayer to which the MAC Merge sublayer connects.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.1 P 29 L 38 # 127
 Jones, Peter Cisco

Comment Type T Comment Status D

There is an editors note regarding TSSI and resolving back to eMAC or pMAC.

I support adding optional parameters to TSSI to indicate eMAC or pMAC, e.g., an optional parameter on TS_TX.indication() and TS_RX.indication().

Need to cross check with 802.1 folks if they cawewere, may need to show up in 802.1Qbu

SuggestedRemedy

Discussion in group and with 802.1 to resolve direction.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

The editor's not should have been removed. We had that discussion in November and decided to use the parameter.

Cl 99 SC 99.1 P 29 L 38 # 99
 Tretter, Albert Siemens AG

Comment Type T Comment Status D

Editor's Note (to be removed prior to publication) - There might be an ambiguity of whether a TimeSync indication was due to an express or preemptable frame. To resolve that, one of the following is needed:

- two TSSIs, one for each MAC,
- an optional parameter indicating express or preemptable added to the TSSI primitives in Clause 90, or
- an additional value such as P_DETECTED should be added for the SFD parameter

SuggestedRemedy

I think at the last Plenary Meeting (11/2014) we already decided together with 802.1 TSN to add an optional parameter to the TSSI primitives in Clause 90. This optional parameter indivates express or preemptable.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Remove the editor's note

Cl 99 SC 99.1.1 P 29 L 45 # 31
 Hajduczenia, Marek Bright House Network

Comment Type ER Comment Status A

Goals are only important during the life of the project and become meaningless once the amendment is approved and merged into the main document.

SuggestedRemedy

Remove 99.1.1. Goals and objectives are documented in project documentation and they belong there.

Response Response Status C

ACCEPT.

Cl 99 SC 99.1.1 P 30 L 36 # 131
 Vahid Tabatabaee Broadcom

Comment Type TR Comment Status A

In the first goal, the minimum latency requirement for express frame is not clearly specified. Is it LIFO or FIFO latency? Is the assumption that every switch should also support cut-through? The minimum number of octets for a fragment can be set to be larger than minimum packet size (i.e. 64B) therefore the latency for express traffic can be larger than 2 times the minimum packet size plus IPG. Even when pre-emption is supported, with 64B minimum fragment size, latency can be larger than the stated value because we have to add the packet processing latency.

Proposed change is to make the latecny goal more general.

SuggestedRemedy

Provide a maximum latency to initiating transmission of an express frame that arrives to an empty queue which is independent of the size of packet in the way of its transmission (i.e. the preemptable packet)

Response Response Status C

ACCEPT IN PRINCIPLE. See comment #31 resolution which deletes the subclause.

Remainder of the response is only relevant if the proposed resolution of 31 isn't accepted.

This is only specifying latency from the MAC Client interface to the MDI. Bridge queues and bridge latency (e.g. packet processing time) is out of scope.

It isn't clear what change would satisfy the commenter as no text is proposed.

Either:

- leave unchanged (because it can provide that latency when the receiver is configured to support minimum fragment size of 64)
- modify to add plus receiver additional fragment size or
- delete goals subclause

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.1.1 P 30 L 36 # 110
 Jones, Peter Cisco

Comment Type E Comment Status D

Bullet a) says "express frame less than two times the minimum packet size plus IPG." which is not true when the peer asks for addFragSize != 0.

SuggestedRemedy

Decide if we need lots of detail here (in which case update point a), or maybe add another goal which explains the ability to increase min fragment size at the cost of increased express frame insert delay.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. See comment #31

Cl 99 SC 99.1.3 P 31 L 35 # 111
 Jones, Peter Cisco

Comment Type E Comment Status D

Seems like we need to add TSSI interface to Figure 99-2—MAC Merge Functional Block Diagram based on Figure 99-1—Relationship of MAC Merge....

SuggestedRemedy

Add the TSSI interface.

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 99 SC 99.1.3 P 31 L 36 # 67
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status R

Figure 99-2 shows a very strange arrangement - in the transmit direction, we have the Transmit Processing function, but in the receive direction, we have Receive Processing function and then Express Filter - separating pMAC and eMAC processing for some strange reason.

SuggestedRemedy

Merge Receive Processing function and Express Filter into a single functional block that is connected to pMAC and eMAC - there is a way to decide whether the given mFrame goes to pMAC or eMAC based on SMD value and no need to keep two separate functional blocks for no reason.

Figure 99-6 and Figure 99-5 need to be then merged into a single State Diagram and then executed within the Receive Processing function. Remove 99.4.6 Express filter - it is already covered in 8023br_1501_hajduczenia_2.pdf - new text for 99.4.5 Receive processing.

Response Response Status C

REJECT.
 What to put in one state machine and what to separate out is a subjective decision.

For the receive function, there is more complexity because checking and handling error cases must be done. In the editor's opinion, the current structure breaks the machines into more understandable pieces.

Also, the existing Receive Processing state diagram is fairly crowded on one page. If they were merged the diagram would still be split into two figures because it can't readably fit on one page.

There is no interaction needed between the two state diagrams so splitting into two is best.

Cl 99 SC 99.2 P 31 L 41 # 53
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

Based on objectives of the project, MAC merge only supports MAC Clients supporting preemption. Anything apart from that is outside the scope of the project.

SuggestedRemedy

Change text in lines 41-41 to read: "This subclause specifies services provided by MAC Merge to a MAC Client supporting preemption."

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.2.1 P 31 L 46 # 54
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
 The statement in lines 46-47 is not accurate. MAC Merge uses MMSI (not MSSSI)

SuggestedRemedy
 Change text to read: "MAC Merge uses MSSSI to request suspension of transmission of preemptable frames from pMAC."

Proposed Response Response Status W
 PROPOSED ACCEPT.

CI 99 SC 99.2.1 P 31 L 46 # 89
 Tretter, Albert Siemens AG

Comment Type E Comment Status D
 Responsibilities of MAC Client using MAC Merge:
 "The MAC Merge can use the MSSSI to request to a ..."

I assume it is ment:
 "The MAC Merge Client can use the MSSSI to request to a ..."

SuggestedRemedy
 Please correct

Proposed Response Response Status W
 PROPOSED ACCEPT.

CI 99 SC 99.2.2 P 31 L 54 # 90
 Tretter, Albert Siemens AG

Comment Type E Comment Status D
 99.2.2 MMSI
 The model used in this service specification is identical to that used in 1.2.2.

=> The reference to 1.2.2 is not resolved.

SuggestedRemedy
 Please correct

Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 This is a reference to 1.2.2 in the base standard.

CI 99 SC 99.2.2 P 32 L 4 # 106
 Tretter, Albert Siemens AG

Comment Type T Comment Status A
 The MM_CTL.request primitive described in this subclause is mandatory.

Why is this service mandatory.

Preemption can be used also without TAS.

SuggestedRemedy
 Discussion needed

Response Response Status C
 ACCEPT IN PRINCIPLE. Discussion was had no change needed.

It was made mandatory because options are bad and should only be used when it is expected that the burden of supporting a feature is high enough to make it worth dealing with an option - having to figure out which implementations support it and which don't and changing usage when the option isn't supported.

It is not burdensome to support, it is likely that all implementations of MAC Merge will support it and because if it is optional, we would need to add support for indicating whether the capability was present (e.g. managed objects).

The MAC Client can use preemption without ever sending the primitive so it places no burden on the Client.

CI 99 SC 99.2.2 P 32 L 5 # 3
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A
 "The MM_CTL.request primitive described in this subclause is mandatory." - this seems to imply this is mandatory

SuggestedRemedy
 Change to read: "MAC Merge shall support the MM_CTL.request primitive.

Response Response Status C
 ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.2.3.1.1 P 32 L 20 # 4
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
"The hold_req parameter can take one of two values, HOLD or RELEASE." - a simple statement will suffice here

SuggestedRemedy
Change to read: "The hold_req parameter takes one of two values: HOLD or RELEASE."

Proposed Response Response Status W
PROPOSED ACCEPT.

Cl 99 SC 99.2.3.1.1 P 32 L 23 # 14
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status R
Wording can be simplified for text in lines 23-27 to avoid circular explanations and unnecessary fluff.
It is not important what state eMAC is, as far as definition of primitives is concerned.

SuggestedRemedy
Change text to read: "The value HOLD suspends transmission from pMAC by preempting a frame currently transmitted by pMAC and allowing transmission from eMAC, regardless of the content of eMAC queues. The value RELEASE terminates preemption and allows transmission by pMAC, regardless of the content of eMAC queues."

Response Response Status C
REJECT.
The suggested remedey text is incorrect in more than one way. MACs have no queues. Release does not terminate preemption - it allows the pMAC to transmit when there is no eMAC frame to transmit.

Preemption still occurs when the last parameter was RELEASE.

Cl 99 SC 99.2.3.1.2 P 32 L 29 # 5
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
Anything not covered explicitly by teh standard is by definition outside the scope - no need to list every single thing.

SuggestedRemedy
Remove 99.2.3.1.2

Proposed Response Response Status W
PROPOSED REJECT.

This is the format used for primitive descriptins. They always have all the subclauses. When one of the subclauses doesn't apply (and that is almost always the case for either when generated or when received) this is what is put in the subclause.

Cl 99 SC 99.2.3.1.3 P 32 L 35 # 162
Regev, Alon Ixia

Comment Type E Comment Status A
The text describing the "Effect of Receipt" of the MM_CTL.request primitive is not clear.

SuggestedRemedy
Change "The receipt of this primitive with the value HOLD will cause MAC Merge to suspend transmission from the pMAC" to "The receipt of this primitive with the value HOLD will cause MAC Merge to not start transmitting frames from the pMAC and to potentially preempt the current frame from the pMAC if preemption is enabled and minimum fragment size requirements are met"

Response Response Status C
ACCEPT IN PRINCIPLE.
"The receipt of this primitive with the value HOLD will cause MAC Merge to:

preempt if preemption is active, a frame from the pMAC is currently being transmitted and minimum fragment size requirements are met and

to not start transmitting frames from the pMAC

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.2.3.1.3 P 32 L 35 # 6
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
Description of the receipt of primitive can be simplified to remove redundant text

SuggestedRemedy
Change text to read: "When the MMSI primitive with the value HOLD is received, MAC Merge suspends transmission from pMAC and allows eMAC to start transmission. When the MMSI primitive with the value RELEASE is received, MAC Merge suspends transmission from eMAC and allows pMAC to resume transmission."

Proposed Response Response Status W
PROPOSED REJECT. The proposed text is incorrect, see comment #14.

Cl 99 SC 99.3 P 32 L 42 # 7
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
This is the first time we mention "preemption capability" - no need to add a new term.

SuggestedRemedy
Changes lines 42 and 43 to read: "When the preemption function is enabled, MAC Merge sends mFrames to the Reconciliation Sublayer (RS). An mFrame contains either a fragment of a preempted preemptable frame or a whole frame."

Proposed Response Response Status W
PROPOSED REJECT. See comment #48

Cl 99 SC 99.3 P 32 L 44 # 112
Jones, Peter Cisco

Comment Type E Comment Status D
Regarding the editors note starting "This project is only addressing MAC Merge operation over full duplex links."

I think that we should remove this, and add a restricting goal in 99.1 that says something like "support only systems that use the preamble and SFD as defined in 3.2.1 and 3.2.2", to automatically move all forms of EFM out of scope.

SuggestedRemedy
Consider proposed remedy of deleting editors note and adding goal to 99.1

Proposed Response Response Status W
PROPOSED ACCEPT IN PRINCIPLE.
The term is "point-to-point link" - Add to the end of the first sentence of 99.1 on a point-to-point link.

For deleting the note see comment #8

Cl 99 SC 99.3 P 32 L 45 # 8
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A
Remove the editorial note - the scope of the project is clear.

SuggestedRemedy
Per comment

Response Response Status C
ACCEPT. The note was inserted to document consensus of a discussion of your proposal.

Cl 99 SC 99.3.1 P 33 L 1 # 1
Boiger, Christian b-plus GmbH

Comment Type T Comment Status A
The definition of the mFrame format contains also changes to the packet format, e.g. 6 byte preamble, SMD.

SuggestedRemedy
Change subclause title to "mFrame and mPacket format" and use "mPacket" instead of mFrame where applicable.

Response Response Status C
ACCEPT IN PRINCIPLE. There is no need for two terms. mPacket is more consistent with 802.3 use of packet and frame.

Replace all instances of mFrame with mPacket.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.3.1 P 33 L 14 # 171
 Regev, Alon Ixia

Comment Type E Comment Status D

We are inconsistent about the minimum data size (some places state this as 60 and some as 64). I believe we have settled on 64 bytes as the minimum data size for any fragment. I think for the mFrame containing a non-initial fragment or frame we should have the data size as >= 64 octets. For the express or initial fragment, the data size should remain >= 60 octets

SuggestedRemedy

For the "mFrame containing a non-initial fragment of a frame" change the data size from >= 60 octets to >= 64 octets.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE.

Our multiple of 8 octet requirement for data fields of non-final mFrames is what complicates this. That is why we have sometimes 60 and sometimes 64.

The frame CRC counts as part of the 64 octet minimum for mFrames. Therefore the minimum here of 60 is correct - any of these frames could be a final mFrame and have a frame CRC.

The check in the state machine for preempting uses 64 because preemption won't be able to happen until 64 bytes and the MCRC have been sent.

There is another comment to remove the multiple of 8 test.

If the multiple of 8 test is retained, add text to 99.3.5 to include the multiple of 8 requirement for non-final mFrames and point out that this means minimum data in a non final mFrame will be 64.

CI 99 SC 99.3.1 P 33 L 25 # 10
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

Text under Figure 99-3 contains a lot of repetitive statements and could be organized in a more clear fashion.

SuggestedRemedy

Change text in lines 25-28 to read: "The format of an mFrame depends on data it carries. An mFrame carrying an express frame (transmitted by eMAC) has the same format as the express frame. An mFrame carrying a complete (non-fragmented) preemptable frame or the initial fragment of a preempted preemptable (transmitted by pMAC) has the SFD octet replaced with the appropriate SMD value, per Table 99-1. An mFrame carrying any of the non-initial fragments of a preempted preemptable (transmitted by pMAC) has the SFD octet replaced with the appropriate SMD value, per Table 99-1, and includes an additional fragment count octet (FRAG_COUNT) following SMD."

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. Countr should be counter

CI 99 SC 99.3.1 P 33 L 3 # 30
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Remove ":" at the end of the sentence. It is not needed.

SuggestedRemedy

Per comment

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.3.1 P 33 L 3 # 9
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R

For simpler referencing, Figure 99-3 should be divided into two figures.

SuggestedRemedy

Figure 99-3 should show the left part of current Figure 99-3, and new Figure 99-4 should show the right part of current Figure 99-3. Remove text under both fragments of the figure. Once that is done, change the text in line 3 to read: "Figure 99-3 shows the format of an mFrame containing an express frame (transmitted by eMAC) or the initial fragment of a preempted frame (transmitted by pMAC). Figure 99-4 shows the format of an mFrame containing any of non-initial fragments of a preempted frame (transmitted by pMAC)."

Response Response Status C

REJECT. It makes it easier for the reader to see the two formats side by side in one figure and there haven't been any cases where it created a referencing issue.

Cl 99 SC 99.3.2 P 33 L 32 # 11
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

Description of preamble would be simpler if it referenced Figure 99-3 separated into two separate figures, per separate comment. Then we do not need to repeat the rest of information - it is already covered under the figure in descriptive text

SuggestedRemedy

Change text to read: "The octets of the preamble field contain 0x55 (10101010). The preamble is 7 octets long in mFrame shown in Figure 99-3 and 6 octets long in mFrame shown in Figure 99-4."

Proposed Response Response Status W

PROPOSED REJECT. Even if the two were in separate figures, it would be useful to the reader to describe when 7 octets applies and when 6 octets applies.

Cl 99 SC 99.3.3 P 33 L 39 # 15
Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status D

Description of the SMD is full of redundant text and unnecessary side track information that adds nothing to definition of the SMD.

SuggestedRemedy

Revise the content of 99.3.3 as follows:

The value of the SMD indicates whether the mFrame contains an express frame, the initial fragment of a preemptable frame, or any of non-initial fragments of a preemptable frame.

All valid SMD values are defined in Table 99-1.

The SMD in an mFrame carrying a complete (non-fragmented) preemptable frame or any of the fragments of a preemptable frame also indicates the frame number. Information about the frame number prevents reassembling an invalid frame if the final mFrame of one preemptable frame and the initial fragment of the next preemptable frame are lost. The frame number is a modulo-4 count.

SMD-S refers to any of the four SMD values in an mFrame carrying the initial fragment of a preemptable frame. SMD-C refers to any of the four SMD values in an mFrame carrying any of the non-initial fragments of a preemptable frame."

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 99 SC 99.3.4 P 34 L 24 # 32
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D

Figure 99-3 shows the field called FRAG_COUNT and in here, we use frag_count

SuggestedRemedy

Use capitalization consistent with Figure 99-3, i.e., FRAG_COUNT in the whole draft.

Proposed Response Response Status W

PROPOSED REJECT. It is common for figures like this in IEEE 802.3 for labels in figures to use all caps. For example, see figure 3-1 Packet format where PREAMBLE, DESTINATION ADDRESS, etc. are in all caps even though in text they are not.

If frag_count was used, preamble and data also should be lower case.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.3.4 P 34 L 26 # 33
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
 in 99.3.3, we defined SMD-C and SMD-S to reference any value of SMD-C and SMD-S.

SuggestedRemedy
 Change all instances of "SMD-Cx" to "SMD-C". Change all instances of "SMD-Sx" to "SMD-S".

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 99 SC 99.3.4 P 34 L 27 # 113
 Jones, Peter Cisco

Comment Type E Comment Status D
 Typo in "The frag_ount is set to zero for the..."

SuggestedRemedy
 "The frag_count is set to zero for the..."

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 99 SC 99.3.4 P 34 L 26 # 34
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
 "The frag_count field is only present in frames with an SMD-Cx." - likely, in mFrames, and not generic MAC frames.

SuggestedRemedy
 Change "The frag_count field is only present in frames with an SMD-Cx." to "The frag_count field is only present in mFrames with an SMD-Cx."

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 99 SC 99.3.4 P 34 L 27 # 16
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D
 Grammar: "an incorrect frames" should be "an incorrect frame"

SuggestedRemedy
 Per comment

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 99 SC 99.3.4 P 34 L 26 # 35
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
 Description of frag_count is a tad chaotic. We start with detail and then go into generic statement. It should be the other way around.

SuggestedRemedy
 Change the content of 99.3.4 to read as follows:
 "The frag_count is a modulo-4 counter that increments for each fragment of the preemptable frame. The frag_count protects against reassembling an incorrect frames if up to 3 frame fragments are lost.
 The frag_count field is only present in mFrames with SMD-C. The frag_count is set to zero at the start of each preemptable frame, and mFrames with SMD-S do not contain the frag_count field.
 The valid values of frag_count field are shown in Table 99-2."

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 99 SC 99.3.5 P 34 L 54 # 172
 Regev, Alon Ixia

Comment Type T Comment Status D
 We are inconsistent about the minimum data size (some places state this as 60 and some as 64). I believe we have settled on 64 bytes as the minimum data size for any fragment. In section 99.3.5 we still state "60" octets. It hink we should change it to 64.

SuggestedRemedy
 change "60 octets" to "64 octets"

Proposed Response Response Status W
 PROPOSED REJECT. 64 octets would be incorrect since a minimum MAC frame has 60 octets of data and 4 octets of CRC. See comment #171 for further explanation of why it is 64 in some places.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.3.6 P 35 L 1 # 36
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status R

Description of CRC is a tad chaotic. We start with detail and then go into generic statement. It should be the other way around.

SuggestedRemedy

Change the content of 99.3.6 to read as follows:

The CRC field contains a cyclic redundancy check (CRC) for the mFrame, calculated from the first octet of the mFrame (DATA field in Figure 99-3, or FRAG_COUNT field in Figure 99-4) to the last octet of the mFrame (end of DATA field).
 The CRC shall be computed following steps a) through d) in 3.2.9 and then XORed with 0x00-00-FF-FF. The XOR operation is performed to differentiate CRC in the mFrame from the CRC in a regular MAC frame. The CRC in the final fragment of a preempted preemptable frame contains the CRC from the original MAC frame with no XOR performed on the value. The CRC in the mFrame is therefore used to indicate whether the given mFrame contains the final fragment of a preempted preemptable frame.

Response Response Status C

REJECT. The description starts by explaining that the CRC field contains a CRC check and an indication of whether this is a final mFrame - that is starting with the general.

The proposed replacement text is incorrect - it is important for correct operation that the CRC computation begins with the first octet of the frame; not the first octet of the mFrame. This ensures that the CRC of a non-final mFrame is always different from the frame CRC computation at that point. If it was calculated over a different byte range than the frame CRC, there could be times when the non-final mFrame computation produced the same result as the frame's CRC and the final mFrame was treated as a non-final mFrame.

As clearly indicated in the current text - the mFrame computation is running over the same bytes with the same process as the frame CRC and the only difference is the value XORed in the last step.

CI 99 SC 99.3.6 P 35 L 8 # 163
 Regev, Alon Ixia

Comment Type E Comment Status D

extra "the" in the sentence "For other mFrames, it contains the an mCRC value calculated on the data octets of the frame from the first octet of the frame to the last octet transmitted in that mFrame"

SuggestedRemedy

Change "For other mFrames, it contains the an mCRC value calculated on the data octets of the frame from the first octet of the frame to the last octet transmitted in that mFrame" to "For other mFrames, it contains an mCRC value calculated on the data octets of the frame from the first octet of the frame to the last octet transmitted in that mFrame"

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 99 SC 99.4.1 P 35 L 23 # 66
 Hajduczenia, Marek Bright House Network

Comment Type TR Comment Status D

The behavior described in 99.4.1 is not part of project objectives and it is not needed to cover the scope of this Task Force.

SuggestedRemedy

Remove 99.4.1 - when preemption is disabled, only one data path should be active, i.e., pMAC, and eMAC should be enabled together with the preemption function. Otherwise, we are adding an extra prioritization point at MAC Merge - something that 802.3 was usually avoiding by transferring all prioritization tasks to 802.1 layers above the MAC.

Proposed Response Response Status W

PROPOSED REJECT. This behavior was agreed on in consultation with 802.1 TSN during development of the project baseline. MAC Merge is just ensuring the prioritization dictated by the MAC Client in choosing between the two interfaces, not inserting a new prioritization. See also Comment # 12

CI 99 SC 99.4.1 P 35 L 23 # 164
 Regev, Alon Ixia

Comment Type E Comment Status D

"MAC Merge transmission behavior" would be clearer as "MAC Merge transmit behavior"

SuggestedRemedy

Change "MAC Merge transmission behavior" to "MAC Merge transmit behavior"

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.4.2 P 25 L 31 # 61
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

Description of the process of detecting preemption across the link is confusing - it starts with details and then has summary information spread across the text without much organization. It should start with high level summary and then go into details and requirements.

SuggestedRemedy

Revise 99.4.2 to read:

The preemption function is enabled on the given PHY in the transmit direction only if it is determined that the link partner supports the preemption function. The process of discovering the support for the preemption function on the link partner relies on the exchange of the Additional Ethernet Capabilities TLV (see 79.3.6) between link partners via LLDP (see IEEE Std 802.1Q for details).

The preemption function shall be enabled for the given PHY only if the link partner announces its support for the preemption function via an Additional Ethernet Capabilities TLV. The preemption function shall be disabled on link failure or if the support for the preemption function on the link partner is undefined. Once the preemption function on the PHY is enabled, the PHY shall periodically send an Additional Ethernet Capabilities TLV. (see 802.1Q for details.)

Response Response Status C

ACCEPT IN PRINCIPLE.

The suggested text is reasonably clear with a few of exceptions.

"or if the support for the preemption function on the link partner is undefined" should be deleted. Once preemption support by the link partner has been discovered through LLDP and preemption enabled, there is no way for preemption support by the link partner to become undefined. The last sentence is incorrect, PHYs do not send TLVs and there is no need to continue sending the TLV once preemption is enabled.

Also that the TLV needs to be sent in an LLDP frame addressed to the Nearest Bridge group address is a requirement and has been omitted from the description.

Disabling on link failure was made a should rather than a shall because IEEE 802.3 doesn't always ensure that something above the reconciliation layer can determine that a link failure has been detected by the PHY. The PHY link status goes to management, not up through the MAC. Some implementations may be able to see link status and act on it.

Replace the text of the subclause with:

The preemption function is enabled in the transmit direction only if it is determined that the link partner supports the preemption function.

The process of discovering the support for the preemption function on the link partner relies on the exchange of the Additional Ethernet Capabilities TLV (see 79.3.6). Since preemption operates between the local device and the link partner, the Additional Ethernet

Capabilities TLV should be sent in an LLDP frame addressed to the Nearest Bridge group address (see IEEE 802.1Q).

The preemption function shall be enabled only if the link partner announces its support for the preemption function via an Additional Ethernet Capabilities TLV. The preemption function should be made inactive on link failure.

CI 99 SC 99.4.2 P 35 L 31 # 56
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Subclause headings DO NOT need "." at the end

SuggestedRemedy

Remove "." at the end of individual headings in the draft - there are a few instances.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. The editor will attempt to find and fix all instances.

CI 99 SC 99.4.2 P 35 L 33 # 165
 Regev, Alon Ixia

Comment Type E Comment Status D

"An node" should be "A node"

SuggestedRemedy

Replace "An node" with "A node"

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 99 SC 99.4.2 P 35 L 34 # 114
 Jones, Peter Cisco

Comment Type E Comment Status D

typo in "An node with a MAC Merge sublayer"

SuggestedRemedy

"A node with a MAC Merge sublayer"

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.4.3 P 35 L 46 # 62
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status A

The content of 99.4.3 focuses (in 802.1 fashion) to much on non-compliant devices. The fact is that we do not know (and do not need to explicitly know) what they do with frames if they are non-compliant.

SuggestedRemedy

Reword the content of 99.4.3 to read:
 "Some link partners may implement PHYs that are not IEEE Std 802.1Q compliant. Such PHYs may handle mFrames in an expected (forward mFrames unchanged) or unexpected fashion (insert SFD into the forwarded mFrame, drop mFrames, etc.).
 To verify the operation of the preemption function across the given link once the preemption function has been enabled, the local PHY shall send a fragmented TBD mFrame. If no response from the link partner is received within TBD seconds, the preemption function on the local PHY shall be disabled."
 Leave the editorial note in place as is.

Response Response Status C

ACCEPT IN PRINCIPLE. The proposed text doesn't make sense. There is no way for PHYs to be not IEEE Std 802.1Q compliant since 802.1Q says nothing about PHY behavior. The existing text accurately describes one of the concerns that makes verification necessary. The proposed text perhaps jumbles that together with another concern - that there may be some PHYs with proprietary features that may alter mFrames. As the commenter suggests, we don't need to describe all the issues that verification may address. The suggested text doesn't seem any less focused on non-compliant behavior. Therefore, make no change to the first paragraph.

The proposed resolution of how to do verification is incomplete and we have other proposals.

See Comment #128 for resolution on how to do verification.

CI 99 SC 99.4.3 P 36 L 1 # 128
 Jones, Peter Cisco

Comment Type T Comment Status A

Re editors note about verifying preemption, can we look at using IEEE 802.1ag Ethernet CFM Loop-back (LB) message for this. Is this a must, or nice to have? If it's a nice to have, we just need to suggest CFM LB.

SuggestedRemedy

Discuss using 802.1 CFM LB message rather than defining something new.

Response Response Status C

ACCEPT IN PRINCIPLE. CFM doesn't seem likely to be implemented on most systems implementing this standard and it isn't clear how it would detect whether preemption can work over a link.

Adopt verification proposal 8023br_1501_thaler_1_v2.pdf and remove editor's note

CI 99 SC 99.4.4 P 36 L 17 # 132
 Vahid Tabatabaee Broadcom

Comment Type TR Comment Status X

It is stated that the minimum fragment size of Tx may need to set larger than 64B based on the link partner's Rx requirements. Supporting configurable minimum fragment size of Tx based on the receiver side requirement adds complexity to the MAC design and required buffering at the pMAC.

Suggestion is to make support for minimum 64B fragment size at the receiver side a requirement. The Tx side min fragment size can be set to multiple of 64B if needed but Tx fragment size of link partners do not need to be symmetric.

SuggestedRemedy

Remove lines 16 to 20.

Proposed Response Response Status W

Discuss
 TX fragment size is not required to be symmetric. The transmitter on each side of the link adjusts to the receiver on the other side of the link.

We decided to let the receiver control the minimum fragment size it is willing to receive and the transmitter on the other side of the link has to adjust to the receiver's requirement. This does shift burden from the receiver to the transmitter.

Do we want to change the decision?

Alternatives:

Make no change

Reduce addFragSize to 2 bits to reduce the burden on the transmitter (maximum minimum non-final fragments = 256 instead of 512)

Remove addFragSize (minimum is always 64)

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.4.4 P 36 L 18 # 115
 Jones, Peter Cisco

Comment Type E Comment Status D

I think we should define minActFragSize as a concept/term up front.
 minActFragSize is (64 x (1+ AddFragSize)) octets.

This will simplify the goals definition for preemption latency to be
 "Provide a latency to initiating transmission of an express frame less than
 (64+IPG+minActFragSize)."

SuggestedRemedy

define as suggested, and make use of the new term though the draft where appropriate.

Proposed Response Response Status W

PROPOSED REJECT.
 This would complicate the early material
 See Comment #132 and 31.

Cl 99 SC 99.4.4 P 36 L 8 # 63
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status X

Subclause 99.4.4 contains a lot of repetitive material, covering the same topics that were
 already explained and defined in previous subclauses

SuggestedRemedy

Reword subclause 99.4.4 per 8023br_1501_hajduczenia_1.pdf (clean and diff versions are
 provided).

Proposed Response Response Status W

PROPOSED REJECT. The proposed text doesn't correctly describe Transmit Processing
 operation. See comment #14.

In order to provide an overview of how something operates and then later provide a
 complete more detailed explanation there will be some repetition.

The proposed text is only slightly shorter than the current text and that is partly due to
 describing a simpler behavior that isn't what was agreed to in the baselines.

Cl 99 SC 99.4.5 P 36 L 38 # 55
 Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

Subclause 99.4.5 contains a lot of repetitive material, covering the same topics that were
 already explained and defined in previous subclauses

SuggestedRemedy

Reword subclause 99.4.5 per 8023br_1501_hajduczenia_2.pdf (clean and diff versions are
 provided).

Proposed Response Response Status W

PROPOSED REJECT.
 In order to provide an overview of how something operates and then later provide a
 complete more detailed explanation there will be some repetition.

The proposed text assumes that the Receive Processing Function and Express Filter
 Function of the current draft are combined into a single function but doesn't provide the
 state machines to do so.

See also comment #67

Cl 99 SC 99.4.5 P 36 L 42 # 167
 Regev, Alon Ixia

Comment Type E Comment Status D

Sentence containing comma followed by that: "When the SMD contains an SMD-S, that
 indicates the initial mFrame of a preemptable frame."

SuggestedRemedy

Proposed rewording: "An SMD containing an SMD-S indicates the initial mFrame of a
 preemptable frame."

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 99 SC 99.4.5 P 37 L 10 # 168
 Regev, Alon Ixia

Comment Type E Comment Status D

Sentence containing comma followed by that: "When the SMD contains an SMD-C, that
 indicates an mFrame that continues the data for a preempted frame."

SuggestedRemedy

Proposed rewording: "An SMD containing an SMD-C indicates an mFrame that continues
 the data for a preempted frame."

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.4.5 P 37 L 19 # 116
 Jones, Peter Cisco
 Comment Type E Comment Status D
 typo in "If any of the checks is does not pass,"
 SuggestedRemedy
 If any of the checks do not pass,
 Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 99 SC 99.4.5 P 37 L 19 # 91
 Tretter, Albert Siemens AG
 Comment Type E Comment Status D
 99.4.5 Receive processing
 Typo: "If any of the checks is does not pass, the mFrame is discarded ..."
 Proposal:
 "If any of the checks does not pass, the mFrame is discarded ..."
 SuggestedRemedy
 please correct
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE. See Comment #116

Cl 99 SC 99.4.5 P 37 L 30 # 97
 Tretter, Albert Siemens AG
 Comment Type T Comment Status D
 Comment to the Editors note.
 I there no solution for this problem, where the new Frame is not discarded?
 How does this behavior affect the robustness of the suggested solution?
 SuggestedRemedy
 This issue has to be discussed
 Proposed Response Response Status W
 PROPOSED ACCEPT IN PRINCIPLE.
 It will be discussed.
 Only a small subset of the set of potential bit errors on a link would cause this discard. It happens when the SMD of the final fragment is corrupted or not received (which requires a bit error in the PHY start delimiter, preamble or SMD) or if the CRC of the final frame is corrupted to look like a valid MCRC (or the packet has errors that make the final CRC equal to the computed MCRC value).
 Since this is a very small subset of the errors that can cause frame loss, this would cause negligible additional frame loss.
 It is very difficult in many implementations to have a frame in process and then switch to receiving a new frame with no time to do so. Normally there is an interpacket gap to finish processing the last frame including any error handling such as recovering buffers where the errored frame was stored. To not discard the new frame, implementations would have to be able to terminate an errored frame and in 6 bytes start receiving the new frame when they normally have 20 bytes (IPG and preamble) between frames.
 Discuss with TSN tomorrow

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.4.6 P 37 L 42 # 68
Tretter, Albert Siemens AG

Comment Type T Comment Status D

99.4.6 Express filter
The sentence
"If an mFrame contains an SMD-E, express filter passes the RS PLS_DATA.indication, PLS_SIGNAL.indication and PLS_DATA.request to the eMAC."

and the sentence

"PLS_SIGNAL.indication is never produced by Express filter since it does not apply to full duplex PHYs"

are somehow contradictory.

SuggestedRemedy

Conflict should be cleared

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. Delete PLS_SIGNAL.indication to produce:

"If an mFrame contains an SMD-E, express filter passes the RS PLS_DATA.indication and PLS_DATA.request to the eMAC."

CI 99 SC 99.4.7.2 P 38 L 21 # 58
Hajduczenia, Marek Bright House Network

Comment Type ER Comment Status R

The format of Constants, Variables, Functions, Counters, and Timers should follow 802.3 format. It would be ideal to use format employed in 802.3bk/802.3av Clause 76.

SuggestedRemedy

If needed, I can provide revised format and definitions, if source Frame file is provided.

Response Response Status C

REJECT. The commenter does not describe what the objection is to the format used. IEEE 802.3 has a great many state machines and different Clauses have somewhat different formats.

The format used here is the same as used in some IEEE 802.3 Clauses (e.g. Clause 49) and does not conflict with any 802.3 editorial conventions or state machine conventions as far as the editor is aware.

CI 99 SC 99.4.7.3 P 38 L 30 # 117
Jones, Peter Cisco

Comment Type E Comment Status D

Add (and use) variable minActFragSize = minFrag+(addFragSize*64) to improve readability.

SuggestedRemedy

Change as above.

Proposed Response Response Status W

PROPOSED REJECT. I thought about doing that but since the variable would only be used once to produce yet another variable (preempt) it would be needless indirection and not kind to the readers.

CI 99 SC 99.4.7.3 P 38 L 43 # 74
Tretter, Albert Siemens AG

Comment Type E Comment Status D

Service primitive "MM_CTL.request":

This is sometimes called MM_CTR.request but mostly MM_CTL.request (Fig. 99-2).

SuggestedRemedy

Should be unified.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. Replace instances of MM_CTR with MM_CTL

CI 99 SC 99.4.7.3 P 38 L 43 # 75
Tretter, Albert Siemens AG

Comment Type E Comment Status D

Variable description "hold"

Boolean variable that is set TRUE when MM_CTR.request is received with a value of HOLD and TRUE when MM_CTR.request is received with the value RELEASE.

In both cases the variable "hold" is set to TRUE

SuggestedRemedy

In case of MM_CTR.request is received with the value RELEASE the variable should be set to FALSE.

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.4.7.3 P 39 L 2 # 118
 Jones, Peter Cisco

Comment Type E Comment Status A

The definition of preempt is a little daunting. I spent a while looking at this, and I would really appreciate a little text explaining it (or maybe a pointer/xref if the text is somewhere else already).

SuggestedRemedy

Consider further explanation of why the preempt definition is like it is.

Response Response Status C

ACCEPT IN PRINCIPLE. We normally don't explain transition conditions. The variable preempt is just a way to condense a transition condition that wouldn't fit in the state machine - something that has been done elsewhere in 802.3 for similar situations.

For the commenter's edification - we only preempt when:
 Preempt is active
 An eMAC frame is ready to start or HOLD is asserted
 The minimum data for a fragment has been sent
 and
 Data for a minimum fragment remain.

No change needed in the spec.

Cl 99 SC 99.4.7.3 P 39 L 6 # 80
 Tretter, Albert Siemens AG

Comment Type E Comment Status D

Description Variable "pTX":
 Boolean variable that is set TRUE when there is an ePLS_DATA.request and FALSE otherwise.

should be changed to:
 Boolean variable that is set TRUE when there is an pPLS_DATA.request and FALSE otherwise.

SuggestedRemedy

Change ePLS_DATA.request to pPLS_DATA.request

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 99 SC 99.4.7.3 P 39 L 7 # 86
 Tretter, Albert Siemens AG

Comment Type E Comment Status D

Variable "pTxCplt": In the Transmit Processing State Diagram Figure 99-4 the variable is written as follows: pTXCplt.

SuggestedRemedy

It should be unified "pTXCplt".

Proposed Response Response Status W

PROPOSED ACCEPT. Make the x lower case

Cl 99 SC 99.4.7.4 P 39 L 45 # 169
 Regev, Alon Ixia

Comment Type E Comment Status D

"minfrag" should be "minFrag"

SuggestedRemedy

change "minfrag" to "minFrag"

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 99 SC 99.4.7.4 P 40 L 16 # 119
 Jones, Peter Cisco

Comment Type E Comment Status D

Typo in definition of RX_MCRC_CK. Precscient function returning a Boolean value

SuggestedRemedy

Precscient function returning a Boolean value

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.4.7.4 P 40 L 18 # 77
Tretter, Albert Siemens AG

Comment Type E Comment Status D
.... received after the next 32 mPLS_DATA.indication "primitives" and the next 32 pPLS_DATA.indications...

SuggestedRemedy
Please correct primitives => primitives

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 99 SC 99.4.7.4 P 40 L 35 # 81
Tretter, Albert Siemens AG

Comment Type E Comment Status D
Description SMDS_ENCODE(frame_cnt) function:
Returns an 8-bit vector with the SMD encoding for an SMD-I with frame count of frame_cnt.

SMD-I should be replaced by SMD-S

SuggestedRemedy
SMD-I should be replaced by SMD-S

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 99 SC 99.4.7.4 P 40 L 8 # 76
Tretter, Albert Siemens AG

Comment Type E Comment Status D
Function: pRX_DV(data_valid):
Produces an ePLS_DATA_VALID.indication.

SuggestedRemedy
In this case it should be "Produces an pPLS_DATA_VALID.indication."

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 99 SC 99.4.7.5 P 40 L 40 # 59
Hajduczenia, Marek Bright House Network

Comment Type ER Comment Status R
There is nothing different about counters and variables.

SuggestedRemedy
Merge Counters into Variables.

Response Response Status C
REJECT. There is something different about counters and variables. I expect to be able to increment counters.

Many other clauses in 802.3 list counters separately.

CI 99 SC 99.4.7.7 P 42 L 10 # 64
Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
"FALE" should be "FALSE" in INIT TX PROC

SuggestedRemedy
Per comment

Proposed Response Response Status W
PROPOSED ACCEPT.

CI 99 SC 99.4.7.7 P 42 L 10 # 83
Tretter, Albert Siemens AG

Comment Type E Comment Status D
The variable "txFrame" used in the statemachine described in Figure 99-4—Transmit Processing State Diagram, state "INIT TX PROC".
Typo: "resumeTx <= FALE"

Should be replaced by:
"resumeTx <= FALSE"

SuggestedRemedy
Should be corrected

Proposed Response Response Status W
PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.4.7.7 P 42 L 14 # 92
Tretter, Albert Siemens AG

Comment Type T Comment Status A

In the Tx state machine (Figure 99-4—Transmit Processing State Diagram) the variable "resumeTx" is set to True in state "PREMPT WAIT", but there is no state besides the initialization state "INIT TX PROC" where this variable is resetted to False.

SuggestedRemedy

This variable should be set to False in state "IDLE TX PROC"

Response Response Status C

ACCEPT.

CI 99 SC 99.4.7.7 P 42 L 16 # 84
Tretter, Albert Siemens AG

Comment Type E Comment Status D

Typo: variable "start ipg_imer" used in the statemachine described in Figure 99-4—Transmit Processing State Diagram, state "IDLE TX PROC".
Typo: "resumeTx <= FALE"

Should be replaced by:
"resumeTx <= FALSE"

SuggestedRemedy

Typo: start ipg_imer => start ipg_timer

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 99 SC 99.4.7.7 P 42 L 29 # 93
Tretter, Albert Siemens AG

Comment Type T Comment Status D

In the Tx state machine (Figure 99-4—Transmit Processing State Diagram) the The variable "txFrame" is not incremented at all.
The variable is initialized in state "INIT TX PROC".

The variable has to be incremented after a preempted frame was sent completely.

This is happen after state "P TX COMPLETE" was processed with "pTXCplt=TRUE"

SuggestedRemedy

Add the the necessary operation (txframe++) as suggested.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. See Comment #83

CI 99 SC 99.4.7.7 P 42 L 3 # 120
Jones, Peter Cisco

Comment Type E Comment Status D

Is there any way to move the "The Receive Processing State" closer to in Figure 99-5.
Same for Express Filter State Diagram & Figure 99-6.

SuggestedRemedy

Locate explanatory text closer to figure if possible.

Proposed Response Response Status W

PROPOSED REJECT. I'm not sure I understand the comment. What are you asking to move closer to the diagrams?
Normally the varibles, counters, etc go before the state diagrams and the function descriptions go before that so if that is what you are asking to move closer together, it isn't possible.

CI 99 SC 99.4.7.7 P 42 L 3 # 166
Regev, Alon Ixia

Comment Type E Comment Status D

"Figure" is reapeated twice in "Figure Figure 99-4"

SuggestedRemedy

Change "Figure Figure 99-4" to "Figure 99-4"

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.4.7.7 P 42 L 34 # 65
 Hajduczenia, Marek Bright House Network

Comment Type T Comment Status D
 "start ipg_timer" - timers are loaded / started in 802.3 in a different way: [start interval_timer, BER_Monitor_Interval]. See 77.1.5 for timer related conventions that make most sense.

SuggestedRemedy
 Copy notations for timers from 77.1.5 into 99.4.7.1 and then make changes in Figure 99-4, 5, and 6 where timers are used locally.

Proposed Response Response Status W
 PROPOSED REJECT. Most of IEEE 802.3 starts timers the way they are started here. It follows conventions established by my task group when editing Clause 14 (14.2.3.2) before 1990.

Clause 77 created an extension to that convention because they had a need to load different starting values into timers, but that extension is not needed here.

Cl 99 SC 99.4.7.7 P 42 L 35 # 85
 Tretter, Albert Siemens AG

Comment Type E Comment Status D
 Correct the name of the timer "ipg_time" to "ipg_timer" in the Tx state machine (Figure 99-4—Transmit Processing State Diagram) at state change from state "RESUME WAIT" to connector "B".

SuggestedRemedy
 Correct typo

Proposed Response Response Status W
 PROPOSED ACCEPT.

Cl 99 SC 99.4.7.7 P 42 L 39 # 98
 Tretter, Albert Siemens AG

Comment Type T Comment Status D
 Figure 99-4—Transmit Processing State Diagram
 State change from "PREMPTABLE TX" to "P TX COMPLETE" =>
 Condition "pTxCplt=TRUE":
 Is it not necessary to add the condition "*preempt=FALSE"?

SuggestedRemedy
 Discussion needed

Proposed Response Response Status W
 PROPOSED REJECT. It is not necessary. Preempt and pTxCplt=TRUE are never true at the same time because Preempt requires that there is at least a minimum mFrame left to transmit and pTxCplt=TRUE means there is nothing left to transmit.

Cl 99 SC 99.4.7.7 P 42 L 42 # 133
 Vahid Tabatabaee Broadcom

Comment Type TR Comment Status D
 The txFrameCnt should be incremented in P TX Complete

SuggestedRemedy
 In P TX COMPLETE one line should be added:
 txFrameCnt ++

Proposed Response Response Status W
 PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.4.7.7 P 42 L 44 # 170
 Regev, Alon Ixia

Comment Type T Comment Status A

"PREMPTABLE TX" transitions to "TX MCRC" and then to "P TX COMPLETE" when preempt=TRUE, and it is assumed that preempt will remain TRUE so that "P TX COMPLETE" will transition to "PREMPT WAIT", but this is not necessarily TRUE. As preempt is controlled also by the state of HOLD, and hold could go away at any time, preempt could become FALSE between the "PREMPTABLE TX" and "P TX COMPLETE" states.

SuggestedRemedy

Split "P TX COMPLETE" into two states: "P TX CMPLT FRAG" and "P TX CMPLT FINAL"

Both "P TX CMPLT FRAG" and "P TX CMPLT FINAL" should contain the function "mTX_CPLT".

"TX MCRC" should transition unconditionally to "P TX CMPLT FRAG"
 "P TX CMPLT FRAG" should transition unconditionally to "PREMPT WAIT"

"PREMPTABLE TX" should transition to "P TX CMPLT FINAL" if pTXCplt=TRUE.
 "P TX CMPLT FINAL" should transition unconditionally to "IDLE TX PROC"

Response Response Status C

ACCEPT IN PRINCIPLE. Good catch. There is a simpler fix. If one is in P TX COMPLETE because of a frame preemption, pTXCplt will be FALSE. So change the transition to pTXCplt=FALSE instead of preempt= TRUE.

Cl 99 SC 99.4.7.7 P 42 L 49 # 96
 Tretter, Albert Siemens AG

Comment Type T Comment Status D

Figure 99.5.Receive Processing State Diagram:
 In state "INCREMENT FRAG CNT" the variable "ntxRxFrag" should be incremented and not "rxFragCnt".

SuggestedRemedy

Change rxFragCnt to ntxRxFrag

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 99 SC 99.4.7.7 P 42 L 5 # 60
 Hajduczenia, Marek Bright House Network

Comment Type ER Comment Status D

States in state diagrams should have a name following the pattern: WORD1_WORD2_WORD3 etc. In this way, when referring to individual states, it is not possible to confuse the name of a state with anything else.

SuggestedRemedy

Change "INIT TX PROC" to "INIT_TX_PROC", "IDLE TX PROC" to "IDLE_TX_PROC", etc. Make changes in Figure 99-4, 5, and 6

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. Reluctantly accept - since they are in all caps and followed by the word state, there isn't much chance of confusing them with something else, but it does seem to be the convention.

Cl 99 SC 99.4.7.7 P 42 L 50 # 2
 Boiger, Christian b-plus GmbH

Comment Type TR Comment Status A

The transmit processing state diagram currently enters the state EXPRESS_TX automatically after preempt is set to true. This means it enters the EXPRESS_TX state, also after a hold request is received. Therefore, it also enters the EXPRESS_TX state, even if there is no frame to transmit. The hold signal is e.g. sent based on the 802.1Qbv schedule. There seem to be many cases where a hold signal is sent but no frame to transmit, e.g. during start up.

SuggestedRemedy

After PREEMPT_WAIT one should only enter EXPRESS_TX if eTX=true. If eTX is not true during the time interval "hold" is true, the next fragment should be sent.

Response Response Status C

ACCEPT IN PRINCIPLE. Add "*" eTX=true" to the transition from PREEMPT_WAIT.

Add a transtion from PREEMPT_WAIT to RESUME_PREAMBLE with the same conditions as the transtion from RESUME_WAIT to that state.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.4.7.7 P 42 L 9 # 82
Tretter, Albert Siemens AG

Comment Type E Comment Status D

The counter "txFrame" used in the statemachine described in Figure 99-4—Transmit Processing State Diagram, state "INIT TX PROC" should be unified. In this state the counter is named "txframeCnt".

SuggestedRemedy

Rename the counter txframeCnt to txframe

Proposed Response Response Status W

PROPOSED ACCEPT.

CI 99 SC 99.4.7.7 P 43 L 29 # 95
Tretter, Albert Siemens AG

Comment Type T Comment Status R

Figure 99-5—Receive Processing State Diagram
The case where the mCRC check at the end of a fragment is wrong because of an error is missing.
It should lead to a discard. In the current state diagram it leads to FRAME COMPLETE.

SuggestedRemedy

Add the necessary check and action to the state machine.

Response Response Status C

REJECT. There is no need for this. If the value is not an mCRC value than it is assumed to be a frame CRC.

If that value isn't the correct frame CRC, the MAC handles the frame discard as usual for a bad CRC.

CI 99 SC 99.4.7.7 P 43 L 29 # 94
Tretter, Albert Siemens AG

Comment Type T Comment Status R

Figure 99-5—Receive Processing State Diagram

Where is the cumulative mCRC computation over received fragments?

It should be in the loop of state "P RECEIVE DATA".

SuggestedRemedy

Add a function in state "P RECEIVE DATA" where the mCRC is processed.

Response Response Status C

REJECT. The mCRC computation is fully defined in text. I don't see what would be added by creating a function and adding it to the state machine.

If we did it, we would have to do the same for transmit for consistency.

CI 99 SC 99.4.7.7 P 43 L 43 # 79
Tretter, Albert Siemens AG

Comment Type E Comment Status D

In Figure 99-5—Receive Processing State Diagram at the exit condition for the loop to state "WAIT FOR DV FALSE" the condition "+ OTHER" should be changed to "+ ERORR" as it is stated in the function description of SMD_DECODE on page 40.

Please change the term "OTHER" also in state machine:
=> Figure 99-6—Express Filter State Diagram => Page 44; Line 19
=> Figure 99-5—Receive Processing State Diagram => Page 43; Line 17

SuggestedRemedy

Please correct

Proposed Response Response Status W

PROPOSED ACCEPT.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

Cl 99 SC 99.4.7.7 P 43 L 44 # 100
Tretter, Albert Siemens AG

Comment Type T Comment Status D

Figure 99-5—Receive Processing State Diagram
Transition from state "CHECK FOR RESUME" to state "CHECK FRAG CNT"

In front of the "CHECK FRAG CNT" the check of the "Frame Number" is missing!

SuggestedRemedy

The check of the frame count should be added.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. See Comment #135

Cl 99 SC 99.4.7.7 P 43 L 45 # 78
Tretter, Albert Siemens AG

Comment Type E Comment Status D

In Figure 99-5—Receive Processing State Diagram the state "CHECK CHECK FRAG CNT" contains two times "CHECK"

SuggestedRemedy

Please delete one "CHECK"

Proposed Response Response Status W

PROPOSED ACCEPT.

Cl 99 SC 99.4.8 P 44 L 36 # 57
Hajduczenia, Marek Bright House Network

Comment Type E Comment Status D

State diagrams should be the last subclause describing the given process.

SuggestedRemedy

Move 99.4.8 to 99.4.7 and then 99.4.7 to 99.4.8.

Proposed Response Response Status W

PROPOSED REJECT. This is the position for delay constraints in other Clauses

Cl 99 SC 99.4.8 P 44 L 48 # 123
Jones, Peter Cisco

Comment Type T Comment Status D

I think we need a formulae/express for HRT that takes into account minActFragSize. When minActFragSize > 0, HRT is greater than 1217 bit times.

SuggestedRemedy

Update definition for HRT to account for addFragSize/minActFragSize.

Proposed Response Response Status W

PROPOSED ACCEPT IN PRINCIPLE. See comment # 144

Cl 99 SC 99.4.8 P 44 L 48 # 144
Mitsuru, Iwaoka Yokogawa Electric Cor

Comment Type T Comment Status D late

Maximum HRT value of 1216 bit times is not satisfied if addFragSize is non-zero .

SuggestedRemedy

Modify the last sentence of the 2nd paragraph of 99.4.8 as follows:
HRT shall be no more than $(1216 + 512 * \text{addFragSize})$ bit times.

Proposed Response Response Status W

PROPOSED ACCEPT.
See also Comment #101 which may change the base value.

IEEE P802.3br D1.0 Ethernet: Interspersing Express Traffic 1st Task Force review comments

CI 99 SC 99.4.8 P 44 L 50 # 101
Tretter, Albert Siemens AG

Comment Type T Comment Status A

[Editor's note (to be removed prior to publication) - 1140 bit times is the transmit time for a preemptable mFrame that is too small to be transmitted plus an IPG: 7 octets preamble plus 2 octets SMD and frag_count plus 119 octets data plus 4 octets CRC plus 12 octets IPG. 1216 allows for 96 bit times to recognize the HOLD and prevent beginning transmission of a preemptable mFrame.

SuggestedRemedy

The numbers of bits and octets in the editors note are not quite clear to me.
For example 1140 Bits are a not an even number of octets (142,5 octets).

Discussion necessary.

Response Response Status C

ACCEPT IN PRINCIPLE.
The calculation needs some corrections as it wasn't updated when changes were made. It should be: 6 octets preamble plus 2 octets SMD and frag_count plus 119 octets data plus 4 octets CRC plus 12 octets IPG.

That produces 1144 bits

HRT will be increased to 1240 to keep the 96 bits of reaction time.

"too small to be transmitted" in the editor's note should be "too small to be preempted"