

Section 2 (MAC) CRG Summary

John Lemon

Steve Wood

David James

Summary

- 77 comments for Clauses 5, 6, 8, G, H, I

Clause	Editorials	Technicals	Total
5	4	9	13
6	6	34	40
8	2	10	12
G	0	2	2
H	1	1	2
I	3	5	8
	16	61	77

- All comments (technical and editorial) addressed
- 3 comments rejected
- 10 comments deferred to WG
- 1 comment (57) punted to LME

Issues Covered - Clause 5

- FCS errors: deferred
- Edge wrapping

Issues Covered - Clause 6

- TTL setting: deferred
- Replication
- State table corrections and enhancements
- STQ sizing
- New shaper text
- Access timers
- Edge wrapping details
- Ringlet versus east/west: deferred
- Delay and jitter bounds: deferred

Issues Covered - Clause 8

- Minimum data frame size (0 length payload)
- Frame format redefinitions: deferred
- Invalid field combinations ignored
- New control type for new topology extended status message

Issues Covered - Annexes G,H,I

- Corrections for new crc16 text
- Update example code to reflect changes in last 2 drafts
- Separate Annex I into 2 annexes: one dealing with vector signals and spatial shaping, and the other dealing with datapath/fairness scenarios.

Rejected Comments

■ 55

- Clause 5.3.1.5, page 53, line 33
- Comment: Update text to state that the client does not get back its own multicast transmissions.

d1.1 text

The MAC does not reflect frames back to the client. If a client receives a frame with a SA value of the local MAC address, it does not cause an MA_DATA.indicate primitive to be sent to the originating client.

- Resolution: Multicast is included in the original description.

Rejected Comments

■ 59

- Clause 5.3.2.4, page 54, line 37
- Comment: The effect of the indicate primitive is that the frame is transferred to the client

d1.1 text

The effect of receipt of this primitive by the MAC client is unspecified.

- Resolution: We can't specify what the client does.

Rejected Comments

■ 113

- Clause 8.2.3, page 110, line 51
- Comment: Need to clarify that the SA is global scope, and may not be the address of the station that puts the frame onto the media. Same scope comment about DA.

d1.1 text

The source address field specifies the station sending the frame. The source address contains an individual 48-bit address as defined in 5.2 of IEEE Std 802-1990.

Comments Deferred To WG

- 70
 - Clause 6.3.2, page 68, line 6
 - Comment: The value of the TTL cannot be a local implementation choice. We should mandate what the TTL should be set to - and then how to use the TTL can be a local decision.
This may be especially true considering the recent Bridging proposals, which require specific values of the TTL.

Comments Deferred To WG

■ 71

■ Clause 6.3.4, page 68, line 23

- Comment: Ringlet selection values should be based on span semantic (east / west) interface rather than ringlet semantic 0/1. MAC should map the span parameter on the MAC service interface to appropriate ringlet encoding on the MA_DATA.request primitive and vice versa on the MA_DATA.indication primitive.

The rationale for this is as follows :

There is a logical problem with ringlet 0/1 parameter through the service interface rather than the span. The problem with using ringlet is in the association of symmetrical traffic flows through the service interface. A symmetrical traffic flow is one where transmit/receive traffic flow through a given span or set of spans on the ring. By having the service parameter selected based on span basis rather than ringlet basis allows symmetrical tx/rx flows to carry the same designation through the MAC interface vs. opposite designations. For example, request.east / indication.east refer to transmit receive flows through the same interface. Whereas request.ringlet0 / indication.ringlet0 refer to transmission on completely different spans!!! Most protocols and transmission models will probably use a symmetric transmission model. Data protocols will tend to TX/RX on a single span, and not on a single ringlet. Using the ringlet construct means symmetrical transmission will transmit on one ringlet and receive on the other. This is non-intuitive, very confusing, and will lead to implementation consistency problems.

Comments Deferred To WG

- 73
 - Clause 6.3.4, page 70, line 35
 - Comment: The Setting of TTL is too flexible to be useful.

Comments Deferred To WG

■ 75

- Clause 6.3.4.7, page 70, line 34
- Comment: typo "destination" should be "hop count"

d1.1

Stations shall set the TTL for data frames to no less than the destination to the desired station and to no more than MAX_STATIONS.

Comments Deferred To WG

■ 78

- Clause 6.4, page 74, line 38
- Comment: 6.4, in defining service classes, implies but never defines the methodology for setting bounds that can be met by traffic within each service class. The MAC is only one component of the ring delay. However, it is the crucial element, and the most subject to delay variation that is not easily computable. Therefore, if our standard is to be useful to our customers, we must provide guidance on how one can calculate the ring performance for each service class in terms of our MAC capabilities and other parameters such as number of stations, ring segment lengths, etc. This information could be a part of or an entire informative annex. However, although it will be only informative, the Standard does not fulfill one of its basic requirements, if that information is omitted.

Comments Deferred To WG

- 86
 - Clause 6.6.2, page 79, line 23
 - Comment: "...at least the ring size" requires a station to know how many nodes are on the ring.

Comments Deferred To WG

- 108
 - Clause 8, page 107, line 1
 - Comment: The existing frame format has problems with respect to support of no frame duplication and reordering as required for transparent bridging.

Comments Deferred To WG

- 109
 - Clause 8.1, page 107, line 42
 - Comment: The frame formats have a multitude of problems, including:
 - 1) The fairness frame is special and hard to process.
 - 2) The header is an oddball size of 18 bytes.
 - 3) Strict ordering/no-duplication for multicasts and bridges is not supported.

Comments Deferred To WG

- 1392
 - D1.0: Clause 8.4.1, page 110, line 27
 - Comment: Using Bit 0 of the ring control field as a parity bit for fairness messages does almost nothing to protect this field. In the other frame formats the bit is reserved and thus set to 0. Parity is supposed to detect all single bit errors. In about 1/4th of the instances, a single bit error in the packet type field of a data frame will convert the frame to a fairness frame with valid parity. The FCS should detect that these are invalid frames but if we are relying on that why do we need a parity bit.

Comments Deferred To WG

- 3469
 - D0.3: Clause 8.5, page 86, line 1
 - Comment: Fairness packets do not have a HEC. This is a problem for Type B fairness messages which are broadcast.

Action Items

- Kshitij Kumar requested to provide a contribution on latency bounds for classB (see D1.0 #209)
- Bob Love and rest of WG requested to provide performance text including delay for all service classes (see #78)?

Open Issues

- More precise jitter bounds for classA and classB
- Ringlet selection for wrapped station
- Frame format issues including:
 - Bridging, misordering, duplication effects on frame format and data path
 - Fairness frame format
- Up to date illustrative code