Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal

Resolutions of Ballot on Draft Standard D4.0

Comments WITH RESPONSES on clauses 10 and 11

Seq #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
1	10.2	jz	t		XXGET.confirm should include a possible status of "attempt to read write-only MIB attribute". XXSET.confirm should include a possible status of "attempt to set a read-only MIB attribute. Perhaps a table of possible result values is in order.		**EDITORIAL / Consistency XXGET and XXSET status value changes inserted into 10.2. ACCEPTED**
2	10.3.1.1	jz	t		"Should include a mandatory enable/disable Boolean" is damn straight. And the details of the parameters should be specified.	Finish writing the section. It obviously is only half done.	**EDITORIAL / Consistency Text in 10.3.1.1 was incomplete and inconsistent in format with susequent subclauses on other primitives. This has been fixed by including parameters which carry values for the only 2 MIB attributes which are related to this service primitive. ACCEPTED**
3	**11 D Annex D**	**vh**	**e**	**n**	**inconsistent attributes in MIB and ASN1 descr**	**aHandshakeoverhead is not defined in MIB (section 11.4.4.2...)**	**(fixed in Annex D changes) EDITORIAL/CONSISTENCY ACCEPTED**
4	**11 D Annex D**	**vh**	**e**	**n**	**inconsistent attributes in MIB and ASN1 descr**	**aRateFactor in ASN-1 descr is not defined in MIB; may have been renamed into aMaxRate**	**(fixed in Annex D changes) EDITORIAL/CONSISTENCY ACCEPTED**
5	**11.1.1.1 11.1.2.1 11.1.1.2**	**sab**	**E**	**n**	**These sections contain virtually the same information. 11.1.2.1 is more specific than 11.1.1.1. Suggest an editorial change to combine but keep meaning**	**Editorial (soory I don't have the text to hand)**	**Requested change is editorial, but unnecessary. DECLINED**

Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
	11.1.2.2				Same is true for 11.1.1.2 and 11.1.2.2		
6	11.1.1.1 11.1.1.2	sab	e	n	Beacons contain timestamp elements containing the value of a STA TSF timer at the reference point defined in 11.1.2.3 - not copies of TSF timers, time elements or any other inconsistent term in these sections	eg in 11.1.1.1 ...the AP shall periodically transmit special frames called Beacons that contain the value of the TSF timer of the STA in the AP at the reference point defined in clause 11.1.2.3 Elsewhere tighten up the language - this is after all a specification, not a story book!	Requested change is editorial, but unnecessary. DECLINED
7	11.1.2 11.1.2.3	sab	e	n	First paragraph in 11.1.2.3 is duplicated from 11.1.2 Combine & have a single definition	editorial	EDITORIAL The redundant paragraph in 11.1.2.3 has been removed, the text in 11.1.2 is more specific and is the proper place. ACCEPTED
8	11.1.2.2	jz	t		It is not clear what happens if something *other than* a beacon starts during the random delay, or if another transmission is taking place. Item 3) should say "if the reception of a beacon has not commenced during the delay period" since the beacon probably will have started but not yet completed when my delay ends.		EDITORIAL / Clarification First part of comment DECLINED as unnecessary: The paragraph is concerned with making a decision based on whether or not a Beacon has arrived. The arrival of any other type of frame is irrelevant to this procedure. Second part of comment ACCEPTED by modification of Item (2) to explicitly extend the delay to the end of a reception which is in progress at the end of the defined time interval.
9	11.1.3	sab	e	n	First sentence of this section is duplicated at the start of the third paragraph - remove one or the other	editorial	EDITORIAL 1st sentence removed.

Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
							ACCEPTED
10	11.1.3	jz	t		Clause 10 defines primitives for starting and stopping scanning. It is not clear how the MIB variables referred to here are related to those primitives.	Harmonize	**The requested change is unnecessary, since the existing text is clear. The definition of the service primitive in 10.3.2.1 refers to the same enumerated type names as are permitted for the MIB attribute in 11.1.3. DECLINED**
11	11.1.3	jz	t		It is not clear to me that the low-order 46 bits of the address are the right ones. I though it was the low-order bit of the first (and thus low-order, right?) octet was the Individual/Group bit. I think it should be high-order 46 bits.	Check 802.1 definition of bit order for addresses, or be more precise about which bits are which.	**EDITORIAL Reworded to refer to the bits using appropriate terminology, without reference to the relative significance of the octets of the 48-bit address. ACCEPTED**
12	11.1.3.2.1	jz	t		The last sentence of the section is a consequence of the immediately preceding sentence since the AP is always the one that sent the most recent beacon in an IBSS (Infrastructure BSS – see my general comment on ABSS/IBSS)	Delete the sentence "If the station is an Access Point…"	**Statement may be redundant, but deletion is unnecessary. DECLINED**
13	11.1.3.2.2	jz	T		This section is so terse that it is technical gobbledygook. Item b) is particularly egregious, since it sounds like you can pounce on the medium the instant it frees up, rather than waiting a DIFS like a good little kid.	<<I will rewrite the section during the meeting if required>>	**EDITORIAL / Clarification The problem is due to obsolete references to CCA when the basic access procedure in clause 9.2 is the correct reference since D3.0. A clarification is needed because the existing text does not deal with the possibility that the reception of probe response in progress when the timeout expires. (Several other clarifications of this type pertaining to timeouts while waiting for frame reception were incorporated at the May, 1996 meeting, this instance was**

Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
							overlooked.) ACCEPTED
14	11.1.3.3	jz	t		Wireless Access Points scare me. It seems like the only way to implement such a beast the way things are in D4.0 is to have two separate stations – one that is associated with another BSS and one that forms the WAP's BSS. The restrictions here and in nearby sections need to be reexamined if anyone thinks that it should be possible to have a WAP that is associated with more than BSS in a time-shared arrangement or whatever…		NO ACTION REQUESTED The commenter appears to have overlooked some of the implementation approaches. However, just because this commenter is scared is no reason to change the standard. NO CHANGE
15	11.1.5	jz	t		To support Bob O'Hara's layer-purity ethic, we could provide a way for the PHY to look at the MAC's TSF timer on the fly and a way for the MAC to ask the PHY how soon the next medium-unavailable event is scheduled to happen. Then this section could be generic and apply to any circumstance where the PHY knows of an impending even that will cause the medium not to be available.	Just a thought…	NO ACTION REQUESTED. This appears to be another attempt to solve a problem resolved at the May, 1996 meeting. However, inadequate detail is provided and the procedures specified in the D4.0 draft are adequate. NO CHANGE
16	11.2.1.4	jz	t		I still hate the sentence "No MSDUs or management frames received for stations operating in the Active mode shall be buffered." There is a confusing failure throughout 11.2 to distinguish between buffering for purposes of implementing power saving and other reasons why frames are buffered.	Sigh!	EDITORIAL / Clarification Stated that buffering is for power management reasons. ACCEPTED
17	11.2.1.6	jz	T		Broadcast/Multicast reliability was not fixed before. I still think it is a problem for the reasons Matt originally asserted, and we simply haven't been able to agree on a fix. Here is yet another idea: If bit 0 in the TIM is set, randomize PS-Poll times over an interval that is at least aCWmin times the number of bits in the TIM that are set. (Or 4× or N× or whatever…) This would give the multidestination frames a fighting chance of not getting clobbered by the PS-Polls. Also, the random delay could be distributed between some non-zero value and the larger value, so the AP gets a chance to start sending first	We should discuss and vote on what to do. I can prepare text to describe whatever mechanism the group is in favor of.	This requests a major functional change in an area where other proposals have been voted down by this group at least twice. This new proposal is neither described in sufficient detail to assess its probable benefits and drawbacks, nor is there general agreement that the referenced problem needs fixing. DECLINED

Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
					and is only waiting a random contention time between them. (Or we could relieve the AP of the backoff after every transmission requirement in this case so it gets to send all the multidestination frames before anyone else starts sending Polls.		
18	11.2.2.1	jz	e		Change "Ad Hoc" to "Autonomous". This jives with my suggested nomenclature with the "I" in IBSS standing for Infrastructure, and calling ad hoc BSSs ABSSs. It still preserves the term ATIM since it starts with an "A".		**EDITORIAL** **Unchanged because base comment was attempting to unify on "independent" terminology, so "autonomous" is as much a dangling reference as "ad-hoc" and the statement is the explanation of the source of an acronym, not a specification of protocol behavior.** **DECLINED**
19	11.2.2.1	jz	t		It is not clear to me when the transmission of multidestination traffic happens. The 6th paragraph makes it sound like you can only send directed traffic, since multicast ATIMs are not acknowledged.		**EDITORIAL / Clarification** **Clarified to indicate that broadcast & multicast ATIMs are not acknowledged** **ACCEPTED**
20	11.2.2.4	jz	t		Item l) is physically impossible as stated. I think it should say "be discarded because of excessive buffering time". Again, the text fails to distinguish between doing something (in this case discarding frames) for the particular reason of excessive buffering time and the general activity (of discarding frames).	The text should be more precise.	**EDITORIAL** **The existing text appears to be clear as written. In the opinion of the processing group this does not need to be changed.** **DECLINED**
21	**11.2.2.4**	**AS**	**t**	**y**	**ATIMs are management frames and thus have sequence numbers. In a senario where a fragment burst is partially completed before a beacon transmission, the beacon is transmitted, an ATIM announcing the remaining portion of the frame is transmitted; the receiving station would lose the sequence number information assoicated with the partially completed MSDU.**		**PART OF A NO VOTE** **Declined because there is no necessity for the requested change, as it is a simple design optimization rather than a fault in the protocol, as sequence numbers in ATIM frames can safely be ignored as the effect of**

Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
					The only solution appears to be that the receiving station either ignore the sequence number in the ATIM or we make the ATIM a control frame. I vote for making the ATIM a control frame.		accepting a duplicate ATIM is null. This response has been read to the commenter by phone and he has agreed (after inclusion of another, editorial, change in clause 9.2.8) to change his NO vote to YES. DECLINED
22	11.3 also 11.1.3.2 .1 also 8.1	mif	T	n	There is nothing specified, either procedurally or in the MAC MIB to define an upper bound on the time between receipt of an Associate or Reassociate request at an AP and the generation of the expected response. This leaves open the possibility of independently–implemented stations and APs, both of which are fully conformant with this standard, but which are NOT INTEROPERABLE! In particular, in the case where the AP never responds to these requests before the STA has ceased listening. For power-managed stations, waiting "forever" is a poor alternative. I strongly recommend that we apply the time limits already in the MIB for aMinProbeResponseTime and aMaxProbeResponseTime to the request/response exchanges for Association, Reassociation, and Authentication (for each step in the authentication sequence), as well as for Probe. There also needs to be a constraint that the AP (or responder in the case of Probes and Authentication sequences in an IBSS) make its first attempt to transmit the response within aMinProbeResponse of receipt of a valid request. The requirement for conformance & interoperability is to have an upper bound on the response time between successful receipt of the request and the first attempt to obtain control of the medium to transmit	Add a sentence to each sub-section which defines when response frames are sent. The general format of this sentence is: "The station shall generate and attempt to transmit a XXX Response frame within aMinProbeResponeTime of receipt of a valid XXX Request frame."	Duplicate of 96/106-3, #1 This requests a small, but significant, technical change. This change may be able to prevent certain interoperability problems, but has been found to be insufficient to be a general solution to the cited problem, because (at least) the same situation can occur with authentication frames, and the proposal only covers Association and Reassociation frames. There may be similar problems with other managment frame exchange sequences, and there may be a more general fix, assuming this problem is adequately in need of fixing to justify a change in the future. DECLINED

Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
					the response. With this time interval known, there is a basis for interoperability that allows local decisions at the stations as to how much longer (if any) to wait due to medium access delays, and whether to retry, look elsewhere, etc.		
23	11.4	jz	e		There are omissions from the MIB. Add:	aHandshakeOverhead	**EDITORIAL / Consistency The Annex D is wrong, and aHandshakeOverhead is being deleted there. DECLINED**
24	11.4.1.1.1 11.4.2.1.1 11.4.4.1. 30	jz	e		aCFAware should be called aCFPollable (old nomenclature).		**EDITORIAL / Consistency This was a global change adopted in March, 1996. This instance was overlooked. Corrected here, may need editorial change in Annex D ACCEPTED**
25	11.4.4.1 .14	jz	t		The default value for aWEPDefault should either be -1 or 256, to allow for future expansion of the number of keys WEP has to choose from.		**EDITORIAL / Consistency Changed to default of 0, as already stated in Annex D. ACCEPTED**
26	**11.4.4.1 .20 & Annex D page 359**	**rn**	**e**	**n**	**The default value mismatch for the attribute aPassiveScanDuration (100 in 11.4.4.1.20 and 50 in annex D page 359)**		**EDITORIAL / Consistency The Annex D has been corrected to match this value DECLINED**
27	**11.4.4.1 .32 11.4.4.2 .21**	**vh**	**e**	**no**	**inconsistent use of units**	**change to kbit/s**	**EDITORIAL ACCEPTED**
28	11.4.4.2 .28	jz	t		PHYDATA.confirm should be PHYTXEND.confirm as the timing reference point. And it is no longer the time to receive the ACK, but the time for the ACK to start being received, so it is really just SIFS + some Rx processing delays		**EDITORIAL / Consistency ACCEPTED Corrected here, may need editorial change in Annex D**

Seq. #	Section number	your ini-tials	Cmnt type E, e, T, t	Part of NO vote	Comment/Rationale	Corrected Text	Disposition/Rebuttal
29	11.4.4.2 .29	jz	t		PHYDATA.confirm should be PHYTXEND.confirm as the timing reference point (see comment on 11.4.4.2.28)		EDITORIAL / Consistency ACCEPTED Corrected here, may need editorial change in Annex D
30	11.4.4.2 .34	AS	t	n	aProbeDelay default value should be aMaxMPDUTime so that the default behavior of scanning does not clobber existing traffic. It can always be changed to 0 to allow a station to ignore existing traffic.		This requests a technical change the necessity of which is unclear. The comment neither substantiates an actual problem with the existing default, nor explains why the altered default will solve said problem. DECLINED
31	11.4.4.2 3 D Annex D	vh	e	n	Incomplete definition?	in aTotalBackoffTime we wander whether we need to include pre- and post-backoff.	No change, the station is either in backoff and this counter applies or not in backoff and this counter is static. Issues of fractional slots are irrelevant because this is an integer. DECLINED