Title: Review of Comments Draft D5.2

Authors:

Bob O'Hara and Phil Belanger

Overview:

We reviewed the comment resolution files to see if the comments were addressed adequately. Overall, we found that disposition of comments was very good. We identified nine comments that we felt needed improvement. The original comments and our recommendations are included in this document. If our recommendations are followed, there is one technical change to the draft that is required.

1. In document 96/156-1R1, the resolution to the unnumbered comment on page two should be changed.

RS Comment

Use of "shall" and PICS: The use of the word "shall" is critically important in IEEE standards. A "shall" mandates a conformance requirement. Therefore, the word should be used SPARINGLY, in precisely those clauses that absolutely require conformance for interoperability or correctness. In addition, EACH AND EVERY "shall" must have an associated entry in the PICS proforma. This has not been done in this standard. The PICS refers generally to sections that contain many shall statements. This in incorrect. There should be a 1:1 correspondence between the number of "shalls" in the document and the number of conformance requirements in the PICS..

Rather than have a lot of "shalls", it is common practice to have a complete detailed description of some desired behavior, either in prose or a formal language/state-machine, then have *ONE* statement, such as: "The MAC shall implement the requirements of the Transmit State Machine as specified in clause x.x.". This allows one PICS entry for a complex entity.

RS Proposal

Eliminate and restructure the use of the term "shall" as indicated, or correct the PICS such that there is a 1:1 correspondence between "shalls" and PICS requirements entries.

Original Disposition

Comment respectfully declined.

It is accepted that there are places in the draft where rather than have a prose description covered by a single 'shall' the text uses 'shall' statements for each of the elements that make up the required function. This is a style issue that does not change the specified functionality. The editing burden of changing the style of the draft at this stage is quite frankly too great to accept this comment at this late stage.

Proposed new resolution:

Accepted, in part. The use of "shall" has been removed from the clauses defining the service interfaces and frame formats. The corresponding entries in the PICS have also been removed. However, the working group feels that the use of "shall" in the remainder of the standard is acceptable as it currently exists. The working group also feels that the PICS is a much more useful item in its current form, as it provides more information to a potential user about the instant implementation. The working group also feels that the PICS contains enough detail when referencing a subclause that the vast majority of potential implementers will receive sufficient guidance to build confirming implementations. Thus, the working group declines the further changes requested by the commenter.

---

2. In document 96/156-5R1, the resolution to comment 18 from RDH needs to be changed.

Comment/Rationale:

Encryption must cover the Integrity Check Value (ICV) as well as the data.

Proposed change:

The top of Figure 35 should be redrawn as follows: (diagram show arrowhead extending to cover the ICV field)

Current Resolution:

Declined.

Having the ICV encrypted would strengthen the WEP. Export restrictions in the WEP design have been checked and the ICV CAN be encrypted(NSA).

Proposed new resolution:

Accepted.

Having the ICV encrypted would strengthen the WEP. Export restrictions in the WEP design have been checked and the ICV *can* be encrypted(NSA).

---

3. In document 96/156-8R2, the resolution to comment 11 from TLP should be changed.

Comment/Rationale:

± 0.0025% is four times the frequency accuracy of most crystals, which are typically ± 0.01% devices. Anything better than ± 0.005% typically requires temperature compensation and consequent power and expense. Is this ± 0.0025% really necessary?

Proposed change:

Resolve the question. A note detailing the rationale for the extra expense of temperature-compensated crystals might be in order.

Current Resolution:

REJECTED:

This accuracy requirement is derived from the PHY specification.

(check with TLP)

Proposed new resolution:

Accepted.

Change the last sentence of clause 11.1.2.4 of draft D6.0 to read:

"The accuracy of the TSF timer shall be ± 0.01%."

Rationale for new resolution:

We think the original resolution was oversimplified.

The PHY may require a crystal with ± 0.0025% accuracy to ensure the accuracy of the oscillators. However, there is no PHY *timing* requirement that is this accurate. Furthermore, this section describes the accuracy required for the MAC's TSF timer. This accuracy is needed to allow STA synchronization in IBSS and to enable power save stations to accurately predict the TBTT so they can check the TIM with minimum power consumption. It makes sense for 802.11 implementers to use the crystal required by the PHY for MAC timing functions. However, when a STA enters the Power save state, it usually turns off the oscillators in the PHY. Therefore, an implementation that strictly adhered to this specification might require two crystals with ± 0.0025% accuracy. Since the MAC does not require that level of accuracy for the desired function, we recommend the change described above.

---

4. In document 96/156-8R2, the resolution to comment 14 from MAF should be improved.

Comment/Rationale:

There is nothing specified, either procedurally or in the MAC MIB to define an upper bound on the response time for Management frames other than Probes. There is a risk that conformant implementations might not be interoperable in the absence of of such a bound on the time before the responding station attempts to send Association Response frames, Reassociation Response frames, and Authentication frames (for the 2nd through last frames of any defined authentication sequence).

The problem could occur in a case where an AP (or other responder STA in the case of Authentication sequences) is implemented in such a manner that it will never respond to one or more of these request types within the time that some STA implementation considers a reasonable maximum waiting time for such a response…

Proposed change:

Clause 11.3.1:

A station shall associate with an Access Point via the following procedure:

a) The station shall transmit an Association Request to an Access Point with which that station is authenticated.
b) If an Association Response frame is received with status value of "successful", the station is now associated with the Access Point.

If the Association Request fails for any reason, the station may scan for a different Access Point with which to attempt association.

<SNIP…>

An Access Point shall operate as follows in order to support the association of stations.

Whenever an Association Request frame is received from a station and the station is authenticated, the Access Point shall transmit an Association Response with a status value as defined in clause 7.3.1.9. The Access Point shall make its initial attempt to transmit the Association Response frame soon enough after receipt of the Association Request frame that a successful transmission attempt will be complete within aMaxProbeResponeTime of the receipt of the request.

<SNIP…>

Current Resolution:

Partially ACCEPTED:

The changes related to the station were accepted and the text addopted.

Responder requirements cannot be met.

(check with MAF).

Proposed new resolution:

Partially ACCEPTED:

The changes related to the station were accepted and the text adopted.

The Access Point changes were not accepted.

In particular, the group felt that it was unreasonable to have the following requirement:

7.3.1.9. The Access Point shall make its initial attempt to transmit the Association Response frame soon enough after receipt of the Association Request frame that a successful transmission attempt will be complete within aMaxProbeResponeTime of the receipt of the request.

The group feels that specifying a timeout on the requesting station is sufficient.

---

5. In document 96/156-8R2, the resolution to comment 42 from WD should be changed.

Comment/Rationale:

This section describes that in the ATIM window also Multicast frames shall be transmitted. This is not correct. The ATIM frame can have a multicast address, to announce multicast frames, but the frame itself should be send outside the ATIM window.

This then is also consistent with item d of section 11.2.2.4.

Proposed change:

The ATIM Window is defined as a specific period of time, defined by aATIMWindow, following a TBTT during which only Beacon or ATIM frames shall be transmitted.

Current Resolution:

What is the desired operation? Ask the group.

Proposed new resolution:

Accepted.

---

6. In document 96/156-8R2, the resolution to comment 59 from MAF should be changed.

Comment/Rationale:

The object groups in 11.4 (oSMT in 11.4.2.1.1, oMAC in 11.4.2.2.1) are defined according to ISO/IEC 10165–2, whereas the Annex D uses SNMP v2. These should be consistent (unless 11.4.2.x is removed due to another comment).

Proposed change:

Use SNMPv2 in 11.4.2.x

Current Resolution:

ACCEPTED:

(check with MAF).

Proposed new resolution:

REJECTED:

This action was not taken.

The group harmonised on the ISO/IEC 10165-2.

---

7. In document 96/156-9R1, the resolution to comment 1 from GMG (and the dependent comments numbered 2, 21, 25, 31, and 37) needs to be changed.

Comment/Rationale:

Currently the entire MIB is specified to be mandatory for Standard Compliance. Since the MIB is not required for interoperability between stations, this is considered far to restrictive. Therefore its support should be optional, which brings this standard more in line with the other 802 standards, none of which define the MIB to be mandatory. The intend of standardizing should be that when a MIB is provided it should use the definitions defined in the optional MIB.

Proposed change:

Make the Status of all items in PC15 Optional.

Current resolution:

Accepted. The management function will be optional, but if implemented it shall be implemented using the MIB as described in the standard.

Proposed new resolution:

Accepted, in part. Due to comment processing in clause 11, the MIB was significantly revised. Much of the MIB has been deleted and other portions have been put into optional packages. However, there are portions of the MIB that are mandatory.

---

8. In document 96/156-9R1, the resolution to comment 23 from MAF needs to be changed.

Comment/Rationale:

The MAC protocol is described solely in English prose, supported by a few diagrams. There is no formal description of the protocol behavior, either as state machines or as procedures in a programming language. This is a major impediment to interoperable implementations of the standard, especially by people who did not participate in the development of the standard. This commenter believes that, by D5.0, there is a great degree of common understanding of the desired MAC behavior among the people who have been active in the MAC group for the past several years, and that the protocol is both implementable and useful. However, there is little chance that a person(especially one for whom English is not their native language) who has not been involved in a recent meeting of the 802.11 MAC group, will interpret all of the text in clauses 8 through 11 in the same manner that the authors of that text, and the voters who approved D5.0, intended. Rather than attempt to catalog incomplete, ambiguous, or potentially conflicting text in the MAC description, this commenter prefers to concentrate on the development of a set of state machines which provide a more precise description of the desired behavior. Some of the areas which are most likely to be misinterpreted include the relationship among the various long–period intervals (beacon interval, contention free repetition rate, dwell time, listen interval); the interaction of indeterminite duration events (such as delivery of a fragmented MSDU when one or more

MPDUs require retransmission) with time boundaries (dwell boundaries, beacons, contention free periods or contention free medium occupancy limits); and the expected behavior at station and access point for power save poll generation and response. (As an example, read clause 9.2.5.2,then try to find all the exceptions and/or modifications to the backoff rules "defined" therein — this is not a particularly bad definition, but if all stations do not implement backoff in an identical manner, the distributed coordination function upon which this entire protocol is based will not operate fairly, and may not operate correctly! A backoff function in a MAC control state machine can provide a single place where all of the relevant backoff behavior, can be clearly defined.)

Proposed change:

Include a precise description of the desired MAC behavior, either as a set of state machines (preferred) or in a procedural language (acceptable but less desirable). The author of this comment will bring to the 802 Plenary meeting in Vancouver a set of state machines which are an attempt to define the MAC behavior informally described in D5.0. These state machines, which will be in submission P802.11/96–132, could be incorporated directly to become the
contents of Annex C. The simplest way to incorporate a formal description of the
MAC protocol is to insert the state machines into the (presently empty)
Annex C
– MAC State Machines and to change this from an informative annex to a normative annex.
This requires far less restructuring of the text in clauses 8 through 11 than placing the state machines in one or more of those clauses. A statement needs to be added early in the document and/or in the introductory paragraphs of each clause which describes MAC operation than the formal definition is the state machines in Annex C, and in the event of a conflict between the text and the state machines the state machines take precedence.

Current resolution:

Proposal to MAC Group: Accept this comment by:
(a) Deleting Annex D
(b) Making Clause 11 (GMDO description of MIB) correct and in agreement
with draft and normative
(c) Restrict GMDO to SNMP-compatible subset of possible data types
(d) Have Clause 11 MIB grouped to agree with the new optionality criteria
in the PICS.

Proposed new resolution:

Accepted.
State machines or other formal description of the protocol will
be included.

Comment: the existing resolution does not have anything at all to do with the commenter's requested change.

---

9. In document 96/156-9r2, the resolution to comment 32 from WD needs to be changed.

Comment/Rationale:

The specification of the ATIM window is inconsistent between the subject sections. Section 11.4.4.1 specifies 4 KusecAnnex D specifies 1000, while the units are not specified. Suggest to specify 4 Kusec, which will suit the DS and FH Phy.

Proposed change:

Update Annex. D accordingly.

Current resolution:

Accepted TEXT_NOT_CHANGED

Proposed new resolution:

Accepted. Due to resolution of other comments, the specified attribute has been deleted from the MIB. It is now a parameter of MAC Management service primitives. It is consistent in all places that it appears.