Author: ***Simon Black, InTalk***

Subject: ***Supplementary Comments on Draft D5***

The following 'comments' came to light after the submission of my sponsor ballot. Most are consistency issues, but there are a few interesting 'cans of worms'. I'm not really sure about the status of these comments given their timing. However, I offer this paper in the best interest of improving the draft and have no particular requirements for these issues being processed, or for a formal response, etc..

Clause 5.5

Disassociation and deauthentication frames are currently specified both to be class 1 frames. Should disassociation be class 2, and only deauthentication class 1? If the receiving STA is in state 1 with respect to the station sending a class 3 frame, then the last paragraph of 5.5 specifies that a deauthentication frame shall be sent; this is also the case if a class 2 frame is received. Disassociation frames need only be sent if the transmitting STA is authenticated but not associated, in which case the receiving STA is in state 2.

Data frames with both ToDS and FromDS set are currently classified as class 3, which would require two APs that are communicating as part of a wireless DS to be in state 3 with respect to each other. Unfortunately there is no mechanism to permit this.

Class 2 frames include "Data - Directed data frames only (FC control bits "To DS and "From DS" both false" this would seem to be an overuse of the word directed - which is used elsewhere in the standard to indicate unicast (as distinct from broadcast). Incidentally, if not all STAs are mutually authenticated in an IBSS a broadcast data frame might elicit a fair few Deauthenication frames.

Clause 8.3.2

The algorithmic description for transmit specified precludes the use of aDefaultWEPKey if any STA with which this STA needs to communicate does not support WEP, since it then becomes impossible to suppress the use of WEP for a subset of stations. Should a distinction be drawn between a station for which no mapping is present in aWEPKeyMapping, and a station for which the mapping is present, but in which WEPOn is set to False?

The algorithmic description for reception does not currently include the semantics of aExcludeUnencrypted.

It is not clear in this description exactly how broadcast data frames with WEP on are handled.

Clause 9.2.5.2

A station is required to perform backoff:

1. After transmitting an MSDU, with another MSDU ready to transmit
2. After transmitting an MPDU with More Fragments = 0

Isn't the former specification redundant (subsumed by the latter)?

Clause 9.2.5.3

Do the Station Short Retry Count and Station Long Retry Count have any useful function? The absolute values of these counters appear not to be significant; while the CW parameter is specified to change whenever SSRC or SLRC are incremented, these events are in turn derived from events on MSDU-specific retry counts.

Clause 9.2.5.6

The first two sentences currently read:

“The following is a description of using RTS/CTS for the first fragment of a fragmented MSDU. RTS/CTS may also be used for fragments if their size warrants it."

The second sentence would appear to be incorrect: the frame exchange sequences (clause 9.7) do not allow this, and the last sentence of the subsequent paragraph of this clause appears to contradict it.

The language of clause 9.2.5.7, and the MIB description of aRTSThreshold, however, would appear to assert that RTS/CTS should be used on fragments other than the first in a burst, since it is specified that a value of zero “shall be used to indicate that all MPDUs shall be delivered with the use of RTS/CTS”. This statement is not accurate anyway since a control frame is an MPDU.

Clause 9.2.6

The last sentence of the first paragraph reads:

“Multicast MSDUs shall be propagated throughout the ESS".

Are Broadcast MSDUs (which are otherwise mentioned here explicitly) deliberately excluded?

Clause 9.2.8

I believe that the current specification of the receive tuple cache has some anomalies (particularly when considered in conjunction with clause 9.5).

Duplicate detection is intended to prevent delivery of duplicate MSDUs or MMPDUs (hereafter MSDUs, for brevity) at the MAC SAP. Within the MAC, suppression of duplicate fragments falls into three cases:

1. a fragment of a previously completed MSDU
2. a fragment of an MSDU currently being reassembled
3. a fragment of a discarded (partial) MSDU

In the first case, a cache entry <RA,Sequence#> is sufficient to reject all duplicate fragments. In the second, <RA,Sequence#> plus the number of the last successfully received fragment defines uniquely the next fragment number to accept. In the last case (where a partial MSDU has been discarded under clause 9.5 (expiry of aReceiveMSDUTimer)), the current specification requires fragments already received to be discarded, even though no indication has been made to higher layers with regard to this MSDU, and successful reception might be possible if the fragments were accepted (assuming the first fragment received here was fragment 0).

I suggest that the cache should thus contain only <RA,Sequence#> pairs, not tuples including fragment number, and should represent MSDUs delivered to higher layers, not MPDUs presented to the MAC layer. Clause 9.5 would also require amendment to specify that duplicates of previously-received fragments of MSDUs currently being received should be discarded.

Clause 9.3.1

The sentence “The longest delay occurs when the current frame is an MSDU which is larger than both aRTSThreshold and aFragmentationThreshold” is misleading, in that aRTSThreshold should be compared with MPDU size, not MSDU size.

Clause 9.3.3.1

In the tenth paragraph, “... broadcasts and multicasts buffered for such stations shall be sent...”: the PC has no knowledge of which stations belong to any given multicast group, so “for such stations” is redundant.

Clause 9.3.3.4

The assertion that the minimum value for aCFPMaxDuration allows one data frame to be transmitted in each direction only holds under the assumption that the initial Beacon frame is not delayed by a busy medium, and that medium occupancy limits do not intervene.

The specification of the maximum value is also ambiguous, mathematically. Should it be:

1. aCFPRate - aMaxMPDUTime + 2.aSIFSTime + 2.aSlotTime + 8.aACKSize; or
2. aCFPRate - (aMaxMPDUTime + 2.aSIFSTime + 2.aSlotTime + 8.aACKSize)

I believe (2) is correct.

Clause 9.3.4.2

The last paragraph implies three states for CF-Pollable stations:

1. Requested to be on polling list;
2. Did not request to be on polling list, but may be added by PC
3. Requested never to be placed on polling list.

The association/reassociation request frames contain only one bit (capability information, bit 3 (clause 7.3.1.4)) to represent this information. It could be taken that the last state is indicated by clearing the CF-Pollable bit; however clause 7.3.1.4 specifies that a CF-Pollable STA shall set this bit. Clause 9.3.3.1 paragraph 13 appears to imply that the AP can send CF-Poll type frames to a non-CF-Pollable STA in any case.

Clause 9.4

I believe MMPDUs can be fragmented (if their size warrants it). Another one of the many places where MSDU really means MSDU, or MMPDU.

The number of octets in a fragment is specified to be “determined by aFragmentationThreshold” (paragraph 3), and to “depend on” the same (paragraph 4). (The first sentence of paragraph 4 does not otherwise make sense, anyway.) This is, at best, an imprecise specification.

If the intention is that all fragments other than the last should be of size equal to aFragmentationThreshold, then the MIB specification should impose a constraint that it should contain an even value.

It is not clear from this clause alone whether the fragment size can be influenced (at first transmission) by other factors (for example, medium occupancy), although these are explicitly excluded from affecting retransmissions.

The specification that the transmit MSDU timer starts “on the attempt to transmit the first fragment of the MSDU” could be clarified since this is not a precise definition of a point in time.

Clause 9.6

Limiting of transmission speed to aMaxRate appears to introduce some potential inconsistencies if arbitrary values are allowed. Should aMaxRate have a minimum value equal to the highest rate in aBSSBasicRateSet? If not, what is the intended use of this parameter?

During the contention-free period, frames of subtype Data{+CF-Poll}+CF-Ack effectively have two recipient stations: the station to which the frame is addressed, and the station which sent the immediately preceding frame. Transmission rate restrictions (paragraph 4) should presumably apply with respect to both stations.

Clause 9.7

The sequences listed are not quite correct with respect to IFS periods after certain CF-Sequences. The separation between multiple CF-Sequences (listed as SIFS) should be a PIFS interval after a CF-Sequence involving a directed transfer from a polled station to another station in the BSS.

The following sequences appear to have been omitted:

1. CF-Poll(no data){+CF-Ack} - Mgmt(dir)
2. CF-Poll(no data){+CF-Ack} - Mgmt(dir) - ACK

Fragment bursts within the CFP are not clear due to the definition of both frag and data(dir).

Clause 11.1.3

The specification of (in effect) the broadcast SSID as a "DesiredSSID containing a value of zero" is ill-defined, and contradicts the definition given in clause 7.3.2.2 (a zero length SSID).

It is not clear whether a station which is scanning in Passive mode can make use of Probe Response frames directed to another station? Clause 11 is somewhat inconsistent on this - compare 11.1.3 and 11.1.4.

Clause 11.1.3.2.1

"... shall be the only station to respond to probes until a Beacon frame is received.". Presumably this should be limited to Beacons with the current BSSID (or possibly SSID).

Clause 11.1.3.3

The statement that a STA which is not an AP "shall scan for the presence of an existing BSS..." conflicts with clause 11.1.3's assertion that a STA with aScanState equal to False which is not a member of a BSS may start its own BSS without scanning for a BSS to join first.

Clause 11.1.4

Should a STA adopt the TSF value from a Probe Response frame from the AP in their BSS, if that frame is directed to another STA? (If the station is a member of a BSS, it is unlikely to be sending Probe Request frames, since it's already synchronised, and receiving Beacon frames.)

The mechanism by which a station "returns to its previous BSS" is not well defined. In order to send a beacon frame, the station needs to be synchronised to the BSS, unless the frame is sent other than after the TBTT. However, in order to send the newly adopted information, the new TSF value is required, so it would appear that two TSF timers are required (or at least accurate knowledge of the relative values). In an FH system, the intervention of dwell periods may further complicate the issue.

If the station chooses to send a beacon, should it remain in the old BSS until it successfully sends a beacon frame, or should it make only one attempt? What should happen if a further beacon is seen, which supersedes the parameters we’re attempting to send? To which BSS does the station belong at this point?

Clause 11.2.1

What is the meaning of aListenInterval set to zero (the default value)? Does this indicate that the station never listens ? The correct default is probably 1, which would indicate that the station is awake for all Beacon frames.

Incidentally, under the DCF (or in the contention period), the frame exchange sequences (clause 9.7) do not provide for the use of RTS/CTS after a PS-Poll. In cases where a station is hidden from the AP, but within range of the station sending the PS-Poll, a collision with a returned data frame is possible; if RTS/CTS were used, this might alleviate the problem. [Note that it would not completely prevent it, since the duration implied by the PS-Poll allows only for an ACK from the AP, and an RTS requires more air time; however, the collision window is made significantly smaller in this case.]

Clause 11.2.1.2

The last sentence “The AP shall identify those stations for which MSDUs are buffered by setting bits in the TIM's virtual bitmap that correspond to the appropiate SIDs.” conflicts with the definition of TIM (clause 7.3.2.1) “Each bit in the traffic-indication virtual bitmap shall correspond to traffic buffered for a specific station within the BSS that the AP *is prepared to deliver* at the time the beacon or probe response frame is transmitted.”

The latter (clause 7.3.2.1) allows the AP to set bits corresponding to a subset of the buffered MSDUs, for example if it knows that the corresponding station will not be awake at beacon time, or if the traffic to be delivered exceeds aCFPMaxDuration. 11.2.1.2 appears to preclude this behaviour.

The reference to probe response frames in 7.3.2.1 is, however, almost certainly erroneous.