Clause 1.3

Author: **Jay Warrior**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
2	16	86	The use of a line to indicate SAP (s) across the interface is also open to misinterpretation. It would be better to indicate the presence of the SAP with something like a set of parentheses e.g ( ).		accepted: add solid blocks or something to lines where SAPs are

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Partially Accep*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
2	21	31	802.11f specifies that operational context is passed from AP to AP as an STA roams. However, the operational context will be lost if the old AP fails or loses its network link	Rewrite 802.11f so that IAPP can recover from a failed AP or a failed link to an AP without losing operational context for an STA	The suggested remedy would require that state for Stations be stored in the network fabric. Further that state storage could not be in a AP since the desire is to retain the state in the event of AP failure. The TG believes that the complexity of the proposed change is beyond the charter of the TG. However, the 2nd failure mode suggested (of losing a link to an alive AP) is something that can reasonably be handled. The TG has added text to recommend that APs monitor the status of their L2 link to the DSM and if it goes down that the AP dissassociate Associated stations and refuse further associations and reassociations until the link is restored.

Author: **Terry L Cole**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
1	32	6	Introduce RADIUS on first use	suggested_remedy = Use Remote Authentication Dial In User Service (RADIUS)	OK Will do - add radius to def list in sec 3 doc then it defined for all uses. Will add spelling out to first usage also.

3	2	5	There is no obvious connection of Mobile IP to one of the references.	Please help the reader understand where to go to find more in Mobile IP, either by adding a reference obviously connected (obvious in the same way DHCP is obvious) or by adding an explicit reference in the text.	Accepted - we will add a parenthetical pointer to MIP doc here

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
3	9	9	I do not understand bullet item "evolution of the IAPP through multiple versions"	Please provide explanation in the text here or other places. [If I missed it, please make it more obvious, I looked for what you might be describing.]	OK - we will remove bullet from here.

Clause 1.3 (figure 1)

Author: **Jay Warrior**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
2	16	85	Figure 1 is ambiguous. What are the dark grey blocks ? What function do they represent. ?		OK - no change requested. None made. The gray is where there is no connection - it only indicates the absence of protocol in that block. The TG added a sentence to say that the gray blocks are where there is no connection between non-gray blocks.

Author: **Terry L Cole**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
2	16	7	I would like you to label the IAPP SAP in this figure	suggested_remedy = Add IAPP SAP and a arrow pointing to the line described in the text which is the IAPP SAP	Ok - we will add the label to identify the SAP

Clause 1.4

Author: **Michael Seals**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

3	25	71	There may be more than three security risks due to inter-AP communications.	Change the text to read, "...present at least three ..."	Ok will make change to sentence - will remove word three and there just happen to be three examples…

Clause 2

Author: Catherine Berger

Comment Type: ***Editorial*** *Vote:* ***Coordination*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
0	0	81	¸ You mention that all the references listed in Clause 2 are subject to revision; however, you do not state that “when a standard is superceded by an approved revision, the revision shall apply.” This sentence needs to be there if you want users to automatically update to the most recent version.		Ok will add suggested sentence except for use of "shall" which we can’t say in Rec practice.

Author: Peter Ecclesine

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
4	14	38	802.11f specifies the use of IPSEC. IPSEC can be used for authentication and key management. However, in 802.11f it is unclear what aspects of IPSEC are being used and for what purpose	Specify what aspects of IPSEC are being used and for what purpose.	Accepted - text has been added that clarifies which portions of the IPSEC family of specifications are used. Specifically the only portion of IPSEC used is ESP and the reference in 2 has been changed to explicitly reflect this. The TG removed unused RFCs from ref lists in 2.0; with this change the list in 2.0 now specifies the IPSEC documents used.

Author: Terry L Cole

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
4	2	10	standard	change to recommended practice or document	OK - will make change

Clause 4

Author: Terry L Cole

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
6	4	11	"There are four service types…" Is this what you mean to say?	I think you mean more explicitly to state that there are four types of service primitives. Please change if so.	Accepted - corrected.

Clause 4 (Fig 2)

Author: **Terry L Cole**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Agreed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
7	2	12	I agree with commenters whom you rejected in letter ballot: MSC charts would be very much more specific that the graph provided here. I would like to show as well a 4th entity in this figure, the MLME.	Primitive Relationships should be expanded by showing a MSC that include the MLME, APME, IAPP, and the other entity (currently called IAPP generated packets). This other entity is perhaps best given another name from the typical architecture of Figure 1, such as UDP/TCP. I am providing sample diagrams for request and terminate, association request, and move request flows, including a variant for each sequence I found described in the text. Please include these after reviewing and making sure they are matching text as you desire.	Accepted: the TG thanks the reviewer for providing MSC diagrams for consideration - this type of effort is really appreciated by a volunteer organization. The TG has reviewed the submitted charts, made some minor modifications as a result of the review and included them in draft 4.1.

Clause 4 (figure 2)

Author: **Jay Warrior**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
7	1	87	Use a formal sequence diagram to indicate causality and temporal sequencing.		Accepted - The TG has enhanced draft 4.1 with expanded MSC diagrams.

Clause 4.1.2

Author: **Jay Warrior**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

Page	Line	ID	Comment	Suggested Remedy	Resolution
8	7	89	There is no linkage in the specification between the shared secret defined here and any parameter in section 5.3. There is also no indication of how the shared secret is distributed.		Accepted: The TG removed from 4.1.2 the shared secret from the param list. The description of the use of the shared secret is in 5.2.

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
8	1	88	There is no definition of types of the service primitive arguments.		this is an abstract interface and not a programming interface - hence the arguments do not have "types" in the sense of the review comment. For example what would the "type" of the "IP address" argument be? The TG declines to make any change (and none was requested).

Clause 4.1.3

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
8	24	8	Are there implied preconditions prior to generating IAPP-INITIATE.request on the TCP/UDP and 802.2 functions?	If so, please add.	Accepted - there are no preconditions that the TG thought of when reviewing the comment to add to the draft.

Clause 4.1.4

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
8	29	13	Are there implied actions upon receipt of IAPP-INITIATE.request between the IAPP function and the TCP/UDP and/or 802.2 functions?	If so, please ad..	Accepted - there are no preconditions that the TG thought of when reviewing the comment to add to the draft.

Clause 4.10.2

Author: **Hugo Pues**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
15	18	68	IAPP-ADD.request seems to be out of context	Change IAPP-ADD.request by IAPP-MOVE.indication	Accepted: ok- wrong name in sentence - corrected

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
15	22	23	Why is there a context block coming from the new AP to the old AP. I believe this is not useful for any purpose described in the RP.	Please delete unless you really mean to have context flowing from the new AP to the old AP.	yes the TG wanted a CB flow from new to OLD - this was requested in prior review rounds - remember that the CB is opaque to TGf so the use of it up to the entity that fills in the CB contents. No draft change was requested except for this comment explanation and so no corresponding draft change was made.

Clause 4.10.4

Author: **Mike Moreton**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
15	36	59	Says that the APME should issue an IAPP-Move.request when denying a move received from another AP. However there is no indication what the Old BSSID field should be set to.	Specify that it should be set to the value of “New BSSID” in the MOVE.indication primitive.	Declined - in 4.8.4 draft 4, page 13, lines 14&15 the value of the "Old AP" is specified. The TG believe that this is what the reviewer referred to as "Old BSSID". Since the document already says how to determine the value, the TG believes that no change to the draft is necessary.

Clause 4.2.2

Author: **Jay Warrior**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

9	2	90	I couldn't find a formal definition of the enumeration defining these return codes. These are needed to complete the interface specification.		The values are enumerated in the text. The values are not mapped to numbers because this is an abstract interface and not a programming interface. Therefore the TG declines to map the values to numbers and no change has been made to the document.

Clause 4.3.3

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
9	29	14	Are there implied preconditions prior to generating IAPP-TERMINATE.request on the TCP/UDP and 802.2 functions?	If so, please add.	Accepted - there are no preconditions that the TG thought of when reviewing the comment to add to the draft.

Clause 4.3.4

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
9	31	15	Are there implied actions upon receipt of IAPP-TERMINATE.request between the IAPP function and the TCP/UDP and/or 802.2 functions?	If so, please add.	Accepted - there are no preconditions that the TG thought of when reviewing the comment to add to the draft.

Clause 4.4.2

Author: **Srinivas Kandala**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
10	5	44	As comment #21 in the submitted WG level ballot comments indicates, there is no need for the Status. The TG has misinterpreted the comment and gave a wrong disposition. The primitive should be retained but it does not need to have a specific status as the returned value of the Status has only one possible value. When the termination is complete, IAPP-TERMINATE.confirm is indicated to the APME.	Delete the Status parameter in the IAPP-TERMINATE.confirm.	Decline - the TG prefers to be more explicit here - and the feeling is that in the future there may be additional values and it would be more difficult to add the parameter later if that becomes the case. All of the service primitives specified consistently have a status field and the TG prefers to retain the consistency.

Clause 4.5.1

Author: **Terry L Cole**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
10	25	16	STAs that do not properly reassociate when moving.	Please make this text more clear. Perhaps you mean STAs that to not reassociate per 802.11 5.4.2.3 and 5.7.3 but rather issue a new association request? If so, please be more specific.	Accepted - TG has improved the sentence along the lines requested.

Clause 4.5.2

Author: **Jay Warrior**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
10	31	91	I know this has been discussed before, but without enforcement of the condition that sequence numbers obey an increasing, possibly non sequential ordering, no practical use of the sequence number can be made. Lack of this restriction effectively negates the ability to use this for the mechanism that it is intended.		The TG recognizes that the seq number mechanism is not perfect, however it is all that 802.11 provides to TGf. To have something better there would have to be a change to the 802.11 protocol, which TGf is not empowered to do. Also for the purpose of resolving the problem of rapid reassociation, the sequence number is adequate to the task. The comment is declined (to the extent that no change was requested). The TG hopes that the explanation provided will help the reviewer understand the reasoning of the TG.

Author: **Srinivas Kandala**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

10	38	45	This has been a subject of several comments in the past at the WG level ballot. As the author of the section correctly notes, the sequence number will be an ambiguous indication of the most recent association. I have reviewed the November' 01 minutes and I find the determination as made by the algorithm presented is still non-deterministic. Using universal time is a much better way. I also do not accept that using universal time increases the complexity substantially. Compared to the complexity added by IPSec the complexity introduced by universal time is negligible.	Delete sequence number at all instances in the draft and replace it with Universal time.	The TG recognizes that the seq number mechanism is not perfect, however it is all that 802.11 provides to TGf. To have something better, there would have to be a change to the 802.11 protocol, which TGf is not empowered to do. It is also important to remember that for the purpose of resolving the problem of rapid reassociation, the sequence number is adequate to the task - as the sequence numbers will change by relatively small amounts (10-20) not large amounts (1000s) - even in the modulo rollover case it is still easy to determine the order. The TG is sorry that you disagree with the TGs position as the TG firmly believes that the mechanism is sufficient as specified.

Clause 4.5.4

Author: **Arnoud Zwemmer**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
15	12	72	The IAPP Multicast Address is not defined.	Define IAPP Multicast address or change multicast to a broadcast.	The multicast address has been applied for; as soon as it is received, the place holder will be changed to the assigned value. The address will be filled in before the draft is submitted to the stanards board.

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
11	7	30	The Layer 2 Update frame mechanism is unreliable and when it fails communications can be disrupted for long periods.	Define at least a heuristic mechanism to solve problem of lost Layer 2 Updates, if not a recovery mechanism.	The reviewer is reminded that L2 is defined to be an unreliable delivery layer. IAPP is designed to support L2 roaming operation and hence the design requirements do not include perfect reliability. Additionally, a "failure" of the L2 update frame is only an issue until the station next sends a packet. The TG thinks that an additional heuristic mechanism is neither needed or appropriate. The comment having been considered, the suggested change is respectfully declined.

Author: **Srinivas Kandala**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Partially Accep*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
11	8	46	I used 4.5.4 only as a reference. This comment applies to several other sections. Why is this standard defining the maintenance of the forwarding tables? There is already a perfectly good standard for doing this and the implementers should use it which allows forwarding beyond the 802.11 devices.	Delete all instances of Layer 2 Update frame, and the update of the forwarding tables. If forwarding of the tables is deemed important, incoporate 802.1D by reference.	TGf is not performing maintenance of forwarding tables; rather it is causing a specific frame to be issued on the link which has that effect. The frame is a trigger to invoke the 802.1D actions. Therefore the TG believes that the action being performed is in fact what the reviewer has requested. Since the doc already does what was requested, the comment was accepted but no change was necessary to the draft to reflect this.

Clause 4.7.2

Author: **Jay Warrior**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
12	27	92	Replace with a mandatory requirement that sequence numbers be in increasing order, preferably but not necessarily sequential. This property needs to be preserved modulo 2^(N) – 1, where N is the number of bits in the sequence number. It is required that N be specified in this section, I suggest N=32.		The TG recognizes that the seq number mechanism is not perfect, however it is all that 802.11 provides to TGf. To have something better, there would have to be a change to the 802.11 protocol, which TGf is not empowered to do. It is also important to remember that for the purpose of resolving the problem of rapid reassociation, the sequence number is adequate to the task - as the sequence numbers will change by relatively small amounts (10-20) not large amounts (1000s) - even in the modulo rollover case it is still easy to determine the order. The TG is sorry that you disagree with the TGs position as the TG firmly believes that the mechanism is sufficient as specified.

Clause 4.7.4

Author: **Mike Moreton**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
12	22	55	As sequence numbers may wrap, it's difficult to determine whether one is "older" than another. Elsewhere in the document this is correctly noted, but not in this section.	Rephrase the paragraph to make clear that the sequence number is only an aid, not the complete determining factor.	accepted - the text pointed out has been copied from 4.5.2 and used as clarification as requested in 4.7.4

Clause 4.8.1

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
12	34	19	The IAPP-MOVe.request is described as causing frames to be sent to the DS that will udpate forwarding tables for the newly reassociated STA. However, this funciton is attributed later to the IAPP-MOVE.confirm function.	Delete this statement as it properly cannot be done until the IAPP-MOVE.confirm step.	Accepted - text has been corrected in draft 4.1

Clause 4.8.2

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
13	18	20	The Timeout parameter is stated to be waiting on two events. Only one event is correct. The timeout has nothing to to with layer 2 update frame.	Delete this statement.	Accepted - the text has been corrected.

Clause 4.8.4

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

13	25	33	802.11f specifies that the "old AP" is identified by the "old BSSID" obtained from an STA's Reassociation Request. However, the STA does not always know the "old BSSID". For example, when the STA is reset it will often "forget" its "old BSSID". When the STA attempts to associate with a series of APs it may not know with which of the APs it succeeded in associating. In these cases, operational context will be lost.	Rewrite 802.11f so that it can maintain operational context when the STA does not know the "old BSSID"	The situation posed by the reviewer is not possible or desirable. If an station is "reset" then by definition of the 802.11 specification, it can not be in an associated state. Only from an associated state can a station roam - i.e. perform reassociation and only in the reassociation action is the concept of "old BSSID" valid. The second example given proposes that a station will not know what AP it associated with - only a non-802.11 compliant implementation could have this problem since: 1) the association action in 802.11 is completed by a frame that positively acks the association frame 2) only a single association is permitted at any instant since both situations described are not possible under compliant operation of an 802.11 station, the issues submitted can not occur. Therefore, the suggested remedy is not necessary and the requested change is declined.

Author: **Srinivas Kandala**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	*Line*	*ID*	*Comment*	*Suggested Remedy*	*Resolution*
13	31	48	Why isnt the new AP informing other APs about the (re)association with this STA? It appears to me that the step 2 in 4.5.4 should be performed here as well, based on the description in clause 5.	Add step 2 in 4.5.4 to 4.8.4.	the behavior in 4.5.4 is there in order to compensate for badly implemented stations that never do reassociation but only do associations. In 4.8.4 we are dealing with reassociations and therefore are talking about correctly implemented stations and so the extra effort was thought not necessary. The additional step requested is therefore declined.
13	31	43	If update of the forwarding tables indeed is going to be maintained, why isnt the Layer 2 update frame sent?	Clarify or delete all references to the update of the forwarding tables.	Accepted - the correction pointed out has been made in draft 4.1 clause 4.9.3, which is the confirm - since one has to wait to send the update frame until after the confirmation from the old AP.

Clause 4.9.2

Author: **Mike Moreton**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
14	11	56	The Old AP field would seem to be a copy of the information passed in the other direction in the MLME-Move.request. It's not clear what the point of this parameter is - it certainly doesn't add to teh understanding, which is the primary aim of the MLME interface.	Remove this parameter.	The param is present so that it is possible to match confirms to requests - the Old Ap is required to do this. There may be multiple outstanding notifies for the same station. The only way to resolve returning move responses is with the Old AP address at this interface (the APME does not have the IP address that matches Old Ap addresses). The requested change was declined.
14	13	57	The new BSSID field would seem to be passing the address of the local AP to the APME. It's difficult to believe it doesn't know who it is already.	Remove the New BSSID field	If there were only a single BSSID possible, then the comment would be correct. However, It is possible to have multiple WM interfaces and hence multiple BSSIDs - in this situation, the param is required. Change request was declined.

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
14	18	21	The NOT_OPERATING condition makes no sense to me. If an IAPP-initiate has not been issued, why would the IAPP receive a REASSOCIATION.indicate primitive? If the IAPP terminate has been issued already, same question applies. There is no need (as the group points out repeatedly in its comments to readers) to specify what happens when you don't follow the RP.	Remove the NOT_OPERATING return value.	accepted: draft 4.1 changed as requested.

Clause 4.9.4

Author: **Mike Moreton**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

12	35	58	Says that the level 2 update frame is sent by the APME. This would seem to contradict 4.8.1, and is inconsistant with ADD where the frame is sent by the IAPP entity.	Remove the text suggesting that the level 2 update frame is sent by the APME.	Accepted - Thanks for pointing out the error. The frame is sent by the IAPP entity, not the APME. The error is actually that the steps should include waiting for the move-notify response. The text has been corrected in 4.9.3 to reflect this.

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
14	35	17	Technical comment = We are making the APME to do the job of causing the Layer 2 update frames to be issues. I believe this should be the job of the IAPP. The APME in the example architecture in Figure 1 does not have access to the UDP/TCP stack.	Move the description of updating the Layer 2 information to the section 4.9.3 as something that happens when the IAPP-MOVE.confirm is generated by the IAPP function with a SUCCESSFUL status.	Accepted - Thanks for pointing out the error. The frame is sent by the IAPP entity, not the APME. The error is actually that the steps should include waiting for the move-notify response. The text has been corrected in 4.9.3 to reflect this.

Clause 5.1 (Figure 3)

Author: **Mike Moreton**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
18	1	61	The direction of step 7 appears to have been reversed.	Reverse it again.	accepted - the arrow has been corrected.

Clause 5.1.2

Author: **Hugo Pues**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
18	1	69	Message 7 in Fig. 3 "7 -> Move-Req" seems to be wrong	Change by "7 <-¸ Move-Notify"	accepted - the arrow has been corrected and the name changed to reflect the use of packet names consistently in the diagram

Author: **Jay Warrior**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
18	2	93	Replace the figure with a formal sequence diagram.		Declined: The TG has not replaced the figure in 5.1.2 because the Tg believes that it has value. However, the TG has added MSC diagrams to draft 4.1

Author: **Mike Moreton**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
17	31	60	Sentance begining “At this point both APs” seems to have got garbled.	Rewrite: “At this point both APs have the shared secret, and it is used to encrypt all further packets for this exchange.”	accepted - text corrected

Author: **Peter Ecclesine**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
18	1	29	The arrow on the Move-Req is pointing in the wrong direction	Fix	accepted - the arrow has been corrected.

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
17	25	35	802.11f's use of IPSEC requires pairwise security associations to be configured and maintained in RADIUS for each AP pair. This is not scalable or manageable.	Remove need for pairwise security associations	The reviewer should be aware that an AP does not have to maintain a full set of pair wise security association with all other APs in the ESS. The security association is only needed to APs to/from which a station roams. This is a significantly smaller set of information that does enable the use of the pair wise security associations to scale. Further the document was written explicitly to allow an Ap implementation to cache and age security associations to enable an AP vendor to tailor a trade off between performance and cost. The TG believes this is a good design balance for the document and the suggested change is declined.

Author: **Pi-Cheng Law**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
18	1	50	In Fig.3, The arrow (a) of STEP 7 should be reversed and "Move-Req" can be changed into "Move-Notify". Because AP2 (New AP) issues a Move-Req to IAPP, not to AP1 and then IAPP sends a Move-notify packet to AP1 (Old AP). This also corresponds to Fig2's descriptions: The Move-notify packet is sent to AP1 (Old AP) and The Move-response packet is sent to AP2 (New AP).		Accepted - the arrow was corrected.

Clause 5.2

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
18	13	18	The Level 1 support seems not to have a RADIUS client included in the architecture. However, the RP seems to require one in other sections and also seems to demand in the protocol addresses and responses from a RADIUS server.	I think the level 1 should perhaps be removed from the document. If it is kept, then the RADIUS descriptions in section 1 and 4 need to be updated to indicate what to do for level 1 support. For example, provide null or zero partakers for the RADIUS related items. Or for example, the RADIUS client might be a dummy that uses local information rather than a RADIUS server to return the required information.	Accepted: the TG desires to retain the Level 1 support and has altered the draft to eliminate the interdependencies that the reviewer pointed out.

Clause 5.3.1

Author: **Arnoud Zwemmer**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Partially Accep*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

23	16	76	This section says: 1) *should* register as part of the ESS. It is unclear whether IAPP registration is *required* (MUST) to be able to use the RADIUS server subsequently for an MAC address to IP address mapping (level 2 support), or that it is preferable (SHOULD) to do so.	Change 'should' to 'shall' or 'must' or clarify whether or not RADIUS registration is required for level 2 support.	Partially accepted - because the IEEE editing rules for RP documents state that we cannot use the word "shall" which is reserved for standards documents. So while the reviewer is correct from an English point of view, we can not make the English conflict with the IEEE editing rules. Re the point in the 2nd line of the comment, the text has been clarified to correct this.

Author: **Mike Moreton**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
19	17	62	Missing "to" between "communications" and "all APs"	See comment.	accepted - corrected

Clause 5.3.1 (Table 1)

Author: **Mike Moreton**

Comment Type: ***Editorial*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
20	1	63	Tables 1-4 contain references to "note 3" which is a placeholder.	It looks like the gap has now been filled in by table 5, so change note 3 to link to table 5.	accepted - this will be corrected as soon as the numbers applied for are received. Update: the numbers were in the draft in the table - the foornote was incorrect.

Clause 5.3.6

Author: **Pi-Cheng Law**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
23	6	51	In "MLME.REASSOCAIATE.request", the first dot should be changed into a dash.		accepted - corrected

Clause 5.3.7

Author: **Pi-Cheng Law**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
23	14	52	In Table5, "See Table 8/7 "of the description of vendor type 5/6 should be "See Table 11/10"		accepted - corrected.

Clause 5.3.7.2

Author: **Mike Moreton**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
24	14	64	It's not clear whether the next term in the series is SHA1(secret || 2nd SHA1) or SHA1(secret|| 1st SHA1 || 2nd SHA1).	Provide a more complete definition of the series. E.g. "Clear result, then repeatedly set result to SHA1(secret || result) until result has enough bits. (If this is the correct definition).	Accepted: text has been improved in D4.1

Clause 5.4

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
26	13	39	Clause 5.4 brushes off security assurance of a context transfer by stating "crypto protection of the information in the context block, should such protection be required, will be the responsibility of the standard defining the format of the info…." While protection of the block itself "may" be able to be defined in a separate standard, the trust model for AP to AP communications must be assured. No such assurances have been provided anywhere in TGf. How is the new AP supposed to believe authorization information by the old AP? If the old AP is compromised, it can pass invalid authorization records to the new AP unless these records are signed by the AS. The AS must act as the trusted 3rd party and sign such authorization records being passed between the APs.	The comment contains the required changes	The comment is concerned over what could happen if "the old AP is compromised". The draft is securing the traffic between trusted entities, where the entities are APs. The trust of APs is established when they pass the authentication phase of joining an ESS. It is presumed that APs remain trusted during their operation. If an AP become evil during operation, the system has much worse problems that those pointed to in this comment. The fear that some component may be compromised in the future can not mandate that a component may not be used. If that criteria were followed, literally nothing could be used since all components "may" be compromised in the sufficiently distant future. The proposed change is declined.

Clause 5.5.1

Author: **Mike Moreton**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
26	38	65	Same problem as with 5.3.7.2 that the SHA1 expansion series isn't fully defined.	Same solution as for 5.3.7.2.	Accepted: text has been improved in D4.1

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Agreed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
26	35	36	Clause 5.5.1 pg 26 paragraph in lines 35-42.This scheme will not work because the AS is acting as the 3rd party providing authorization between the two APs and yet no 3-party exchange involving the AS is provided. There is no liveness proof in how the new AP gets the old APs security block, nor is there a liveness proof for the old AP to assure that the AS was the one to deliver its security block. How is this to be achieved? Where are the details? Who is making the access control decision, the new AP or the AS?	Provide a prove that security is not breached by the current mechanisms or replace the current mechanisms	Accepted: Draft 4..1 has been changed to always check the timestamp in the ticket that gets delivered to Ap 1 in figure 3. Text was added in 5.5.1 to clarify this.
26	35	40	Clause 5.5.1 pg 26 paragraph in lines 35-42. The expansion function is insufficiently specified and edianness is not clear enough to ensure interoperability between APs.	Provide more details and diagrams to ensure interoperability.	Accepted: the expansion has been corrected in draft 4.1; the endianness was specified in draft 4.1; also SHA1 was changed to hmac-sha1 after consultation with the reviewer and TG.

Author: **Srinivas Kandala**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
26	27	49	It appears to imply that (actually it doesn't say it, but if it doesn't imply what I am going to write, then it has no place in this paragraph) the layer 2 update frame is sent along with MOVE-notify and MOVE-response (probably to clear the entry, perhaps!) to update the forwarding tables. However, sub clauses 4.8.4 and 4.11.4 do not mention the transmission of the Layer 2 Update frame. Something is amiss.	Two options: 1) Delete all references to Layer 2 or 2) Add information about sending Layer 2 Update frame in subclauses 4.8.4 and 4.11.4.	accepted - 4.8.4 corrected; 4.9.3 was corrected to take care of the issue. Thanks for spotting this.

Clause 5.7

Author: **Mike Moreton**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
28	5	66	Says that the ADD-notify is sent to the subnet-local braodcast address. However 4.5.4 says that an IP multicast address is used instead.	IP multicast is probably a better solution, but should really have some configuration option to set it.	Accepted - the actual mechanism is the Multicast, the sentence pointed out was a hold over from prior drafts - the text has been corrected in D4.1

Clause 6.1.5

Author: **Pi-Cheng Law**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
29	13	53	In Table 7, Send-security Block and ACK-security-Block use the General IAPP Packet format. This sentence, the ¡KData field is described in 6.2, 6.4, 6.5 for¡Ktypes, should be added 6.6 and 6.7.		Accepted - draft 4.1 contains the correction

Clause 6.4

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

30	13	41	802.11f uses TCP for handoffs between APs, which has a high cost in terms of set-up, tear-down and maintenance overheads	Remove any mention of TCP	The use of TCP in this instance was seriously considered by the TG during the writing of the document. In the instance where it is used, the TG faced the choice of either inventing a new mechanism to reliably exchange the information required or to use an existing mechanism. In the spirit of a RP, the TG decided to use the existing mechanism of TCP as it was well suited to the required task. TCP does have additional overhead compared to UDP etc - but the overhead is a direct result of the functionality provided. The TG declines to remove the use of TCP since the resulting work would simply replace a well known mechanism with an new, potentially inferior, special purpose mechanism. Removal of the use of TCP was declined.

Clause 6.5

Author: **Hugo Pues**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
31	12	70	... status field are shown in ...	... Table 8.	accepted - corrected

Clause 6.6

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
31	22	37	Remove reliance on RADIUS and/or redesign architecture so that fast and secure roaming is possible.	Add the messages indicated in the comment	The suggested remedy is declined. The TG does not desire to remove all reliance on RADIUS and the comment does not suggest a viable technical alternative. Re the desire for fast and secure handoff, the reviewer is referred to comment #4 from the sponsor ballot and the response to that comment. There may be an opportunity to accomplish the reviewer's desire for fast handoff. The reviewer is encouraged to collaborate with the author of comment #4 to see if they could work further together.

Clause	6.8.16

Author: **Pi-Cheng Law**

Comment Type: ***Editorial*** *Vote:* ***Approve*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
36	10	54	Which one is the length of New BSSID IP address? 4/8 octets in page 36 are different from 4/16 octets described in Table 9.	see the comment.	The correct length is 4/16 the draft has been corrected.

Clause	General

Author: **Arnoud Zwemmer**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
0	0	73	There is too much overhead (registration, using RADIUS) to just obtain a simple MAC-IP address mapping.	Use Inverse ARP to obtain the IP address of the old AP. It is recognized that the DSM MAC address may not be the same as the WM MAC address. However, an AP probably needs to listen promiscusouly on its IP/Ethernet interface anyway, because it must recognize frames not destined for its own address (namely for all associated wireless stations).	Declined: the suggestion to RARP is not acceptable because APs are not constrained to be on the same sub-net.
0	0	74	It is not clear what backend support is needed in an IAPP-aware RADIUS server. The RADIUS message with the standard service type Call-Check seems to suggest a standard RADIUS server is configured with MAC addresses as Usernames and configured to return a Framed-IP-Address attribute. To just allow these MAC Address users access without further authentication seems to open security holes in a RADIUS server that is also used for real strong authentication using 802.1X/EAP-TLS. It is also unclear how this would work with a standard RADIUS server like IAS in Windows. Would MAC addresses need to be configured as users in Active Directory?	Clarify what TGf expects of a RADIUS server, what the exact backend functionality is, whether a standard RADIUS server can be used or that additional backend functionality is required.	Extensions to RADIUS servers are a common occurance when functionality not envisioned during the original development of RADIUS is added to equipment requiring authentication. Many extensions to RADIUS have been created and RADIUS servers provide ways to add additional extensions. The TG disagrees with the suggested remedy and declines to rewrite the draft to use an (undefined) "off the shelf" radius server. It is anticipated that TGf radius extensions will be offered to add TGf functionality to existing server installations - at least one TGf member is planning to do so commercially. Re the potnetial for a security issue mentioned; the access is not via MAC address only, but via MAC address and shared secret.

Page	Line	ID	Comment	Suggested Remedy	Resolution
0	0	75	IAPP must contain a forward roaming facility to facilitate seamless roaming, which is currently missing. Forward roaming allows the current AP to forward state to a potential new AP, so that when the station roams, this state will be already in place at the new AP. Especially in a polled environment, where the AP will only start polling after the station has been added to the polling list, this mechanism will avoid a service interruption. Forward roaming can use similar messages as currently specified for backward roaming (i.e. IAPP-MOVE.xxx), with a few changes. Triggering an IAPP-FORWARD.request requires a message similar to the reassociation request to be added to the MAC. It is recognized that this specific trigger is outside the scope of TGf, but this could be added in TGe.	A) Change MOVE into FETCH. B) Introduce four new clauses for: - IAPP-FORWARD.request { MAC Address; Sequence Number; New AP; Context Blob } - IAPP-FORWARD.confirm { MAC Address; Status, Admission Status } - IAPP-FORWARD.indication { MAC Address; AP Address; Context Blob } - IAPP-FORWARD.response { MAC Address; AP Address; Status} These clauses are essentially copies of 4.8 - 4.11, with a few exceptions 1) 'Old AP' is replaced with 'New AP' 2) Admission Status is included in the .confirm message C) Introduce two new clauses for FORWARD-RESPONSE and FORWARD-NOTIFY packets, which reflect these new messages.	The suggested remedy is declined primarily for the reason that the reviewer noted in the comment: that to implement this functionality there would have to be a change in the operation of the 802.11 protocol and such a change is not within the scope of TGf. However, the reviewer is referred to comment #4 from the sponsor ballot and the response to that comment. There may be an opportunity to accomplish the reviewer's desire for fast handoff without needing to alter the 802.11 MAC protocol. The reviewer is encouraged to collaborate with the author of comment #4 to see if they could work further together.

Author: **Catherine Berger**

Comment Type: ***Editorial*** *Vote:* ***Coordination*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Agreed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
0	0	84	I have completed a SCC10 review of P802.11F/D4 and I find that it meets all the requirements of SCC10 Coordination.		No response required - was a coordination affirmation re format etc.
0	0	83	¸ At the time of submission to the Board, or just prior to publication, you will need to supply a mailing address for each member of the working group that worked on the document. This will ensure that all members of the working group receive a complimentary copy of the standard.		The TG/WG will provide the required list prior to publication.
0	0	82	¸ Please make sure all figures have the appropriate permissions and identifications if any have been taken from another source.		accepted - however, no figures were taken from other sources - all are original to this document. No changes required.

Author: **David Bagby**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

1	Insufficient MIB control – add “AP Management MIB” functionality Issue: When TGf started it’s work, the group decided to support an AP specific SNMP MIB, and a very minimal MIB exists in the TGf draft sent for review. This reviewer believes that the MIB proposed does not provide enough functionality to enable even minimal management of APs. Further it is this reviewer’s position that the ability to mix multi-vendor APs requires more than just the Specific IAPP messages between APs, it also requires a minimal set of common AP management MIB definitions.	Requested Changes to resolve vote: To change this reviewers vote to approve, I request that the committee provide at least the following additional AP MIB functionality as part of TGf: a) Known stations set inquiry The ability needs to be added to inquire and get back a list of stations known by an AP. The current (essentially association) status of each station in the set should be returned with the station set list. The minimum obvious station states would be: Authenticated but not associated; associated (currently active), disassociated (was here, but not here as of when you asked), and re-associated (AP once knew of the station, but it is has re-associated elsewhere). It is not my intent that APs keep history for all time, rather the concept is that AP info about client stations is probably aged and that this inquiry would simply return info about the “known” Stations as of the time of the inquiry. The purpose of the “status” of the Stations being returned with the list is to be able to use the list as input to additional queries. This forms the basis for the ability to use the information in the response as a parameter for additional MIB inquiries that allow one to inquire about information specific to a station or set of stations known by the AP. Returning the Station status provides the easy ability to inquire about arbitrary mixes of stations. For example “asking about currently associated stations” (Stations active with the AP) or “asking about stations that have gone” (common for diagnostic purposes) - or any mix thereof. b) Known Station Attribute Inquiry It needs to be possible to ask about both a single station and an arbitrary set of stations (I suggest a set approach, where for a single station, one simply specifies a set consisting of a single station), and for all stations in the set requested, to get back information that the AP knows about the station(s). Thus, the conceptual parameters of the inquiry are (station set, attribute set that you want to know about). The “station set” input parameter should either be, or be trivially derivable from, the information returned from the “Known Stations inquiry”. I suggest a set approach so that the data consistency issues associated multiple MIB calls over time can be avoided. While the “set approach” is conceptually	Accepted: the TG has expanded the MIB definitions in Draft 4.1.

what is desired, I understand that the ability of SNMP MIB variable definitions may mandate a different approach – this reviewer would consider alternate approaches as specified by TGf.

An additional requirement is that this ability be created in a general enough manner that it can serve as an expandable mechanism as additional attributes for stations are invented. The reviewer requests this as there are multiple active TGs in 802.11 that are inventing additional Station attributes beyond those defined in 802.11-1999. The desire is for TGf to provide the basic framework for asking about per station attributes independent of changes in the attribute sets available.

For the first version of TGf, I want to see the ability to get back at a minimum all the currently defined STA attributes that an AP would know about “its stations”. If TGf believes that additional attributes would be valuable, this reviewer is not opposed to considering enhanced functionality beyond the minimum called out in this review comment.

There are several commercially available sets of MIB extensions within existing products that would more than satisfy this reviewer’s comments. Perhaps TGf could avoid having to invent the MIB details from scratch by soliciting proposals from existing AP vendors.

c) AP operational state control
At a minimum TGf needs to provide the MIB definitions necessary to

1) Deactivate an AP (this essentially requires the ability to tell the AP to disassociate all current stations and not accept new associations).
2) Reactivate an AP (reverse the state above by starting to accept associations again)
3) Reset AP
4) Selectively direct AP to Disassociate a specific Client (one-time event)
5) Selectively allow/disallow a specific STA to Associate/Re-associate to the AP.

The reviewer urges TGf to also consider other AP control abilities (for example those already implemented in many AP’s MIBs).

d) MIB revision level
In order to make the “known station attribute inquiry” in b) expandable, it will be necessary to provide a way

to determine the set of Station attributes known by the AP MIB. This can be accomplished several ways (ex: a separate MIB version call, or the ability to specify the attributes desired via a mask of some type (if attributes are specified that are not known then they ignored in the request). My approval of the draft is conditional primarily on the mechanism being defined being appropriately extensible, not on any specific approach. I believe that the TGf group expertise is best suited to define the details of an appropriate mechanism.

e) AP Identity
It needs to be possible to inquire about basic manufacturer information from an AP. The minimal set includes:
1) Manufacturer ID
2) Model number
3) Revision levels
At first pass TGf may react with “Aren’t these already available in the System MIB for the station?” This reviewer is drawing a distinction between the information that is bound to an AP and the information that is bound to a STA that is conceptually inside an AP. The distinction is important, as architecturally, an AP is an interface between the DSM and the WM, across which it provides DS services.

What is desired is the ability to get version information from the AP entity. That information may well be different than the same info for the AP’s WM STA (of which there may be two in the case of WDS). In fact, recent product approaches have moved the industry toward a place where this will be the likely case as the AP’s STA component is highly likely to change independent of the AP entity itself.

f) AP knowledge about ESS
It needs to be possible to ask an AP what it knows about the ESS that it is a member of. Minimal requirements include:
1) Getting back what ESS is the AP a member of.
2) Getting a list of other APs in the ESS that the AP knows of.

This is intended as a crude way to learn the AP members of an ESS. It would obviously be preferable to “ask an ESS”, however, an ESS is not an entity that one can ask questions of - and inventing such an animal would appear beyond the scope of the TGf work. This inquiry would at least allow some external entity to attempt to build up the set of APs in a ESS.

		f) AP Management MIB access control Clearly, not all the information that is potentially available via the mechanisms above should be made available to anyone who asks. I suggest that TGF specify that the AP Management MIB be restricted to access by other authenticated APs in within the ESS. All that is required is that each AP ignores AP management MIB requests that are from anyplace other than another (same ESS) authenticated AP. This would allow vendors to create an “AP management entity” - which would appear to the other APs as simply another authenticated AP (that probably happens not to accept associations).	
2	BSS ID security Block issues The following comment was received from Justin McCann and I agreed to submit it as Chair of TGf for TG review since he is not part of the Sponsor pool and it does appear to be a significant problem that the TG should address. This problem applies to the New-BSSID-Security-Block in 5.3.7.2, and also the Old-BSSID-Security Block defined in 5.3.7.3 and 6.6. (Section number relative to draft 3.1). The problem as I see it is, once the Security-Block is decrypted, to my understanding there is no way to verify that the decrypted contents are valid. All you have is a bunch of random bytes and no way to verify that they are the bunch of bytes that you want, and that make sense.	Requested Changes to resolve vote: It is my opinion that in order to be able to verify the plaintext contents of the encrypted blocks, you will need to send along separate Radius VSA's that are checksums over the plaintext contents of each block. This needs to be corrected before the draft can be approved.	After much discussion the TG has concluded that the problem presented in the comment is not really the problem it seems to be. For the New-BSSID-Security-Block the contents are protected by RADIUS authentication. For the Old-BSSID-Security block the contents are protected by element ID 14 (HMAC authentication block) as described in 6.6 table 9. Therefore the problem will not occur and no change to the draft was needed. The suggested remedy was declined.

3	Implied static configuration of APs: The current TGF draft calls for significant interaction of APs with Radius servers. Many Access Points will provide RADIUS Client support for authentication services. RADIUS Clients have some well-defined security configuration requirements that will present challenges to effective WLAN deployments. In particular, the RADIUS server must have the Radius Client Shared Secret bound to the Client's IP Address. This essentially requires the Radius Client (the AP) to have a fixed IP address, which is not a DHCP assigned address. That implies the constraint that APs implementing TGf must have statically assigned IP addresses. This reviewer finds that constraint unacceptable. Virtually all the current AP products come out of the box configured for dynamic addressing. It will not be acceptable to MIS managers to have to configure APs with static addresses as part of an installation. This reviewer cannot approve a recommended practice that in requires that all APs be configured with static addresses.	Requested Changes to resolve vote: What is needed is a way for the Radius Client to get the secret into the Client securely. There is a proposal by Robert Moskowitz and John Volbrecht to resolve this issue. TGf needs to recommend that all TGf APs (as Radius Clients) use that proposal (or something functionally equivalent) and that the corresponding Radius server recommended by TGf also support the functionality. A draft copy of the proposed solution has been submitted as a binary attachment with this comment with the permission of the author.	Accepted: The issue in the comment would be a concern for any ESS with more than a few APs; however the need to statically configure the IP addresses can be avoided via the use of IETF draft-moskowitz-radius-client-kickstart-00.txt which can be found on the IETF site. Text has been added to Draft 4.1 to explain this and point to the IETF draft.

Author: **Mike Moreton**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	*Line*	*ID*	*Comment*	*Suggested Remedy*	*Resolution*

0	0	67	If “Broad Market Potential” must be established before work on a project can commence, it is only sensible to check whether that potential still exists before issuing the document. Issuing documents that are of no use to anyone just confuses users. In this case events have overtaken the standard. Interoperability between different vendor’s APs is ensured by the WECA tests - there is no need for an IEEE best practice to do the same thing. While such roaming may be based on associate frames rather than reassociate frame, this is a distinction that is entirely invisible to the user. Secondly this standard provides some additional authentication between APs. This is completely useless as so many other authentication and security holes remain that papering over a few cracks will make no appreciable difference. Finally, there is an apparently sensible context transfer mechanism. However, no 802.11 draft uses this mechanism, so finalising it before even a single use has been identified is premature.	This document should be put “on-hold” until a use that is identifiable to an end-user or network administrator is identified.	Declined - the reason being that the comment is non-responsive per the ballot rules.

Author: **Peter Ecclesine**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Author Emailed***

Page	*Line*	*ID*	*Comment*	*Suggested Remedy*	*Resolution*
0	0	27	802.11f specifies the use of RADIUS as its security infrastructure. However, it is unacceptable to base 802.11f on RADIUS only because:* any future compromise of RADIUS will also compromise 802.11f* RADIUS may not be flexible enough in the future	Rewrite 802.11f so that it can be configured to use any security infrastructure (within reasonable limits), not just RADIUS	The logic provided in the comment does not persuade the TG of the reviewer's position. The fear that some component may be compromised in the future can not mandate that a component may not be used. If that criteria were followed, literally nothing could be used since all components "may" be compromised in the future. If radius was not used and instead "any security infrastructure" was used, then the document would not be a recommended practice - it is necessary to recommend some practice - and in the case of TGf radius is the recommendation.

0	0	34	The reliance on RADIUS means that roaming in many situations will be too slow	Remove reliance on RADIUS and/or redesign architecture so that fast and secure roaming is possible.	The suggested remedy is declined wrt to removing all reliance on RADIUS especially considering that the comment does not suggest a viable technical alternative. Re the desire for fast and secure handoff, the reviewer is referred to comment #4 from the sponsor ballot and the response to that comment. There may be an opportunity to accomplish the reviewer's desire for fast handoff. The reviewer is encouraged to collaborate with the author of comment #4 to see if they could work further together.
0	0	42	It is not possible to understand or validate the 802.11f draft in its current form in any reasonable time because the text is very "dense" and appears not to cover all cases. The draft cannot be passed in its current state because we have no confidence that it does anything useful in a secure manner or that two independent implementers will have any chance of building an interoperable implementations	Provide many more diagrams, probably using some form of formalised state machine, and matching descriptive text. Alternatively, postpone 802.11f until the market matures to the point where it better understands the requirements and appropriate mechanisms for an IAPP	The TG has added MSC diagrams to draft 4.1 which are belived to significantly improve the ability to understand the flow upon first reading. The TG hopes that this will improve the reviewer's general complaint. Other portions of the comment are declined as "non-responsive" under the 802 operating rules. The comment simply states that the reviewer has no confidence but fails to provide sufficient information for the TG to reasonably determine what would be required on a technical basis to satisfy the reviewer. The suggested remedy has two parts; 1) to "add more" - which is also too vague to meet the requirements of a ballot technical comment; and 2) to postpone the publication of the document. Re 2), the TGf project was duly proposed, and authorized and multiple years have been invested in getting to it current draft. The draft presented for Sponsor ballot has passed on the first ballot by 87% and the majority of reviewers approved without any comments.
0	0	28	802.11f specifies a number of extensions to RADIUS. Therefore, 802.11f cannot be used with a standard "of the shelf" RADIUS server	Rewrite 802.11f so that it can use a standard "off the shelf" RADIUS server	Extensions to RADIUS servers are a common occurance when functionality not envisioned during the original development of RADIUS is added to equipment requiring authentication. Many extensions to RADIUS have been created and RADIUS servers provide ways to add additional extensions. The TG disagrees with the suggested remedy and declines to rewrite the draft to use an (undefined) "off the shelf" radius server. It is anticipated that TGf radius extensions will be offered to add TGf functionality to existing server installations - at least one TGf member is planning to do so commercially.

0	0	32	802.11f does not specify how "context" is identified, relying on other standards to specify "context". However, it is not yet clear that 802.11f functionality is suitable for use by any other group.	Identify and liase with other standard groups that are likely to use 802.11f to ensure it is likely to provide suitable functionality.	The comment is declined as "non-responsive" under the 802 operating rules. The comment simply asks the TG to "identify and liaise". In fact, there were joint sessions with other 802.11 Task Groups during the development of TGf and those sessions resulted in the securing of the inter-AP messages, and discussions of the usefulness of the context transfer mechanism. The draft was passed from the WG to Sponsor ballot process and the WG membership is the superset of the 802.11 TG members. It is not reasonable to request an indefinite liaison period with an unspecified set of other "standard groups". It is pointed out to the reviewer that the TGf project was duly proposed, and authorized and multiple years have been invested in getting to the current draft. The draft presented for Sponsor ballot has passed on the first ballot by 87% and the majority of reviewers approved without any comments. Therefore the TG concludes that this reviewers position, which appears to the TG to simply be a tactic for delay, is in the small minority. As there is no way for the TG to reasonably determine what would satisfy the reviewer on a technical basis, the comment was voted "non-responsive" by the TG, making the comment invalid.

Author: **Srinivas Kandala**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Partially Accep*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution
0	0	47	The proposed recommended practice is causing enough confusion due to the discrepancies in clause 4 and 5(and seeing from other comments as well) that it is justified for calling a more formal specification through state machines.	Incorporate state machines.	partially accepted - while state machines are not required for a RP document, we have enhanced draft 4.1 with expanded MSC diagrams which the TG thinks will satisfy the comment. The reviewer is requested to see draft 4.1.

Author: **Terry L Cole**

Comment Type: ***Technical*** *Vote:* ***Disapprove*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	Line	ID	Comment	Suggested Remedy	Resolution

0	0	26	The timeout for the move protocol may leave the STA not associated anywhere. The STA will not be associated at the new AP because the IAPP-MOVE was not SUCCESSFUL. If the timeout occurred while the reques to send IAPP move-notify was in fly, the message may be received correctly by the old AP and the STA disassociated there as well.	Discuss and correct if relevant. If not relevant, add a comment explainig this to the appropriate text.	Accepted - the situation described can occur - however there is nothing IAPP can do about this since the actual reassociation action is taking place via the 802.11 protocol. If the AP causes an 802.11 reassoc response to be sent with the status code "assocaiton denied due to inability to confirm that association exists" (code # 11 table 19 in 802.11) because the IAPP exchanges failed, the Station will have to establish a new Association. An explanation was not added to the IAPP draft as the TGf draft start to become a tutorial on 802.11.
0	0	25	The timeout for the add protocol, may leave the new AP having no idea if the STA has been disassociated from any old AP. Does this matter?	Discuss and correct if relevant. If not relevant, add a comment explainig this to the appropriate text.	Accepted: nope, does not matter because the station gets what it should have expected since it attempting to use an association for reassociation or this may be the first time and there was no prior association. The TG did not feel it was necessary to explain this in the TGf draft.
0	0	22	The terminate protocol does not advise as to any necessary RADIUS related acitivty. However, I believe the threat models described indicate that it is useful to keep status in RADUIS updated. This could keep an imposter from assuming the identify of an AP that is termiating.	Discuss and correct if relevant. If not relevant, add a comment explainig this to the appropriate text.	Declined: There is no info specific to an AP that needs to be disposed of since the initiate does not create individual AP info. However, the TG did not feel it necessary to explain this in the draft.
0	0	24	When is the diasociation expected to occur during a move protocol at the old AP? I cannot find this mentioned.	I have provided a place where I think the protocol for disassociation should be placed in the MSC diagrams I am attaching. Please include descriptions suitable.	Accepted: this should have been in 4.10.4 - draft 4.1 has been corrected.

Author: **William Arbaugh**

Comment Type: ***Technical*** *Vote:* ***ADVISORY O*** *Comment Status:* ***Declined*** *Cmntr Response:* ***Agreed***

Page	*Line*	*ID*	*Comment*	*Suggested Remedy*	*Resolution*

4	We have implemented the current IAPP draft and conducted several measurements of the resultant implementation. The comments included here are based on the results of these measurements, and our desire to support fast and secure roaming between AP's on the LAN and eventually across LAN's. Before we comment directly on the current IAPP draft, we'd like to introduce some background material. Our goals are to allow for fast and secure roaming such that synchronous IP connections such as voice over IP applications (VoIP) will not experience excessive jitter during hand-offs. Current guidelines from the ITU allow for a jitter of approximately 50ms in VoIP connections . This means that the latency from hand-offs of both layer 2 and layer 3 should not exceed 50ms to maintain a quality connection. Before beginning our implementation, we measured the latency of layer 2 hand-off's between commonly available commercial equipment . In this study, we found that current hand-off times far exceed 50 ms. The overall cost, however, was due to the problem of identifying the next AP. This problem, unrelated to IAPP, can be solved independently of IAPP. The main purpose of measuring the layer 2 latency was to establish a base line upon which to compare our implementation of IAPP—determining the total cost of IAPP. We found that the cost of IAPP, as currently specified, using an un-optimized implementation to be approximately 300 ms (NOTE: We believe that an optimized version will reduce this time by one half, but this value (150 ms) is still far too excessive). The main contributor to the cost of IAPP is the reactive nature of the protocol, i.e. the context for the STA is not transferred until AFTER a REASSOCIATION REQUEST message is received by the new AP, and a REASSOCIATION RESPONSE can not be sent until after IAPP completes. As a result, a network utilizing the current IAPP draft will NEVER be able to complete hand-offs quick enough to avoid excessive jitter in synchronous connections and applications such as VoIP and streaming media will suffer significantly.	suggested_remedy = We further believe that the latency problem described, above, can be easily corrected through the addition of one new message type. Our specific proposal will be presented at the next meeting and will included implementation figures which drastically reduce the cost of IAPP within the bounds of the ITU recommendation.	The TG feels that given the advisory nature of this comment and a concern over the delay that adding this functionality at this point in the process would incur (given the lack of available draft text to implement the concepts). It is pointed out that this response will be circulated with draft 4.1, and if there is support from other ballot pool members for this proposal, and the reviewer were to create the text necessary to include the functionality in the TGf draft , that there is potential for inclusion as part of a re-circ comment submission. The reviewer is advised that this would require a completed text proposal by the end of the recirc ballot period which is anticipated for mid December 2002.

Clause	General (Title Page)

Author: **Catherine Berger**

Comment Type: ***Editorial*** *Vote:* ***Coordination*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	*Line*	*ID*	*Comment*	*Suggested Remedy*	*Resolution*
0	0	80	c)¸ The wording of the copyright statement has been changed slightly. Please update it with the following: Copyright © <current year> by the Institute of Electrical and Electronics Engineers, Inc. Three Park Avenue New York, New York 10016-5997, USA All rights reserved.This document is an unapproved draft of a proposed IEEE Standard. As such, this document is subject to change. USE AT YOUR OWN RISK! Because this is an unapproved draft, this document must not be utilized for any conformance/compliance purposes. Permission is hereby granted for IEEE Standards Committee participants to reproduce this document for purposes of IEEE standardization activities only. Prior to submitting this document to another standards development organization for standardization activities, permission must first be obtained from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department. Other entities seeking permission to reproduce this document, in whole or in part, must obtain permission from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department. IEEE Standards Activities Department Standards Licensing and Contracts 445 Hoes Lane, P.O. Box 1331 Piscataway, NJ 08855-1331, USA		accepted - correction made.

1	0	78	Please change the designation from “IEEE Std 802.11F/D4” to “IEEE P802.11F/D4.” The “P” indicates that it is still a draft and not an approved standard. It was done correctly on every page except for the title page.	accepted - corrected

Clause General (Title)

Author: **Catherine Berger**

Comment Type: ***Editorial*** *Vote:* ***Coordination*** *Comment Status:* ***Accepted*** *Cmntr Response:* ***Author Emailed***

Page	*Line*	*ID*	*Comment*	*Suggested Remedy*	*Resolution*
0	0	79	I was a little thrown off by the numbering of this standard. That, in combination with the phrase “Recommended Practice to IEEE Std 802.11, 1999 Edition), made me think this was an amendment, but it is actually a stand-alone document that complements IEEE Std 802.11, correct? To avoid confusion, I would delete the phrase under the designation.		accepted - done.