

Doc 11-03-009R2-F-TGf-Recirc 1 Ballot Comments (by clause,

Clause	1.3
---------------	-----

Author: **Peter Ecclesine**

Comment Type: Editorial Vote: Disapprove Comment Status: Accepted Cmnr Response: Agreed

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
3	23	99	RC1: "looses" should be "loses"	RC1: fix	Comment Accepted:

Comment Type: Technical Vote: Disapprove Comment Status: Accepted Cmnr Response: Agreed

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
3	23	+02	RC1: "The text states that a AP ""should"" essentially cease operations when it loses its ""link"" to the DSM, where the DSM is defined as, ""The medium or set of media used by a distribution system (DS) for communications between access points (APs) and portals of an extended service set (ESS)."" It does not make sense to lose a link to the DSM because the DSM is a ""set of media""	RC1: suggested_remedy = Replace DSM with DS	Comment Accepted: Change DSM with DS

Clause	1.4
---------------	-----

Author: **Bob O'Hara**

Comment Type: Editorial Vote: Approve Comment Status: Accepted Cmnr Response: Agreed

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
3	33	96	RC1: IPsec is no longer used as a generic term in this document.	suggested_remedy = Replace IPsec with ESP, as is done earlier in this clause.	Comment Accept:

Author: **William Arbaugh**

Comment Type: Technical Vote: ADVISORY O Comment Status: Partially Acce Cmnr Response: Agreed

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
3	30	98	RC1: A forged ADD-notify can cause a disassociation for an associated station. The cumulative effect of this is a potential network wide DoS. suggested_remedy = There are two possibilities. 1. Require IPsec for the ADD-notify 2. Be very clear in explaining the potential down side of NOT using IPsec with ADD-notify. line 30 with #2 A bogus MOVE might cause an AP to drop all state it has with a STA, and a bogus ADD-Notify can result in the STA being disassociated. Thus, an attacker with the ability to send IP datagrams to AP's in the ESS can perform a denial of service attack against known STA's. As a result, it is recommended that IPsec be used with ADD-Notify.		RC1: Comment Partially Accepted: Change "...can use IAPP as a Denial-of-Service (DoS)..." to "...can use IAPP or forged management frames as a Denial-of-Service (DoS)..." Move: Butch 2nd Bob M. Vote: Unanimous

Clause	4.10.4
---------------	--------

Author: **Mike Moreton**

Comment Type: Technical Vote: Disapprove Comment Status: Declined Cmnr Response: Open

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
-------------	-------------	-----------	----------------	-------------------------	-------------------

15 36 59

RC1:
There's a requirement to send an IAPP-Move.request primitive. If you go to the description of this primitive, it says that the "Old AP" parameter (sorry I said "Old BSSID" in my original comment) should be set from a field in the MAC reassociation frame. The issue is that in this case the AP never received a MAC reassociation frame – all it got was an IAPP Move-Notify packet. Hence you need to define what this parameter should be set to.

SB comment:
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.

RC1:
SB remedy:
Specify that it should be set to the value of "New BSSID" in the MOVE.indication primitive.

Notes from Discussion: The TG reviewed the sequence diagrams and noted that the commentor had pointed out what was thought to be a subtlety, but in following the description of the proposed problem, it was determined that there is sufficient description already in the draft.

RC1:
Comment Declined: Please refer to lines 14 and 15 of page 13, it indicates that the AP needs to keep the information from the reassociation request (whenever that was) and so the AP does have the "Old AP" value to use when it sends its IAPP-Move.request primitive.

SB resolution:
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	4.5.4
---------------	-------

Author: **Peter Ecclesine**

Comment Type: Technical Vote: Disapprove Comment Status: Declined Cmnr Response: Open

<i>Page Line ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
---------------------	----------------	-------------------------	-------------------

11 7 30

RC1:
"In the last Sponsor Ballot it was noted, ""The Layer 2 Update frame mechanism is unreliable and when it fails communications can be disrupted for long periods."". A request was made to, ""Define at least a heuristic mechanism to solve problem of lost Layer 2 Updates, if not a recovery mechanism"". This request was rejected.

The commenter still believes this in an important issue but understands the previously suggested change might open a whole new can of worms "

SB comment:
The Layer 2 Update frame mechanism is unreliable and when it fails communications can be disrupted for long periods.

RC1:
suggested_remedy = In a spirit of compromise, instead of defining a heuristic algorithm, change, "The IAPP entity sends a Layer 2 Update frame to the DS ..." to, "The IAPP entity sends one or more Layer 2 Update frames to the DS ..."

SB remedy:
Define at least a heuristic mechanism to solve problem of lost Layer 2 Updates, if not a recovery mechanism.

RC1:
Comment Declined: The additional L2 Frames will not guarantee an improved response. Thus adding the suggestion, while helping in very specific cases, can actually hurt in other cases. The TG feels that while the tendered suggested remedy may help in certain cases, it doesn't see a benefit sufficient to accept.

SB resolution:
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.

Clause 4.7.4

Author: **Mike Moreton**

Comment Type: Editorial Vote: Disapprove Comment Status: Accepted Cmnr Response: Open

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
12	22	55	RC 1: You accepted the comment, but the document doesn't seem to have changed. SB comment: 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.	RC 1: SB comment: Rephrase the paragraph to make clear that the sequence number is only an aid, not the complete determining factor.	RC 1: Comment accepted, Sorry to have missed it in the editing process. SB comment: accepted - the text pointed out has been copied from 4.5.2 and used as clarification as requested in 4.7.4

Clause 5.1.2

Author: **Peter Ecclesine**

Comment Type: Technical Vote: Disapprove Comment Status: Accepted Cmnr Response: Agreed

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
-------------	-------------	-----------	----------------	-------------------------	-------------------

17 25 35

RC1:
"In the last Sponsor Ballot it was noted that the 802.11f draft's use of IPSEC requires pairwise security associations to be configured and maintained in RADIUS for each AP pair and that this is not scalable or manageable. I asked that the need for pairwise security associations be removed.

TGf responded, ""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.""

The response highlighted the practical issue related to the configuration of the Radius server. The process of determining which AP pairs need pairwise security associations is likely to be difficult to manage (ie not scalable), particularly as APs are added and deleted from the network and radio conditions change."

SB comment:
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.

RC1:
suggested_remedy = Provide informative text that describes the envisaged RADIUS configuration process

SB remedy:
Remove need for pairwise security associations

RC1: Comment Accepted:
RADIUS doesn't maintain configurable data for the specific AP pairwise groupings. RADIUS acts more like a third party that facilitates the APs to set up inner-communications. RADIUS dynamically generates the keys that are needed. The RADIUS Server tracks the secret of each RADIUS Client, and then the RADIUS Server will provide the security blob that can be used to talk with another RADIUS Client. We changed "The Security Blocks each contain a shared secret for AP-AP connection" to "The Security Blocks each contain information for securing the AP-AP connection. This information is dynamically generated by the RADIUS server as the Security Blocks are constructed. The Security Blocks are encrypted using the AP's BSSID user password (see 5.3.7.2 and 5.3.7.3) in the RADIUS registry."
Also change the "shared secret and it is used" to "information" last line 5.1.2.
Also Change on page 30 "AP's supported ESP and AH transforms," to "AP's supported ESP transforms and ESP authentication algorithms"

SB resolution:
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.

Clause 5.3.1 (Table 1)

Author: **Mike Moreton**

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

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

20 1 63 RC 1:
 Note from chair: Reviewer Accepted SB comment resolution and provided following comment in recirc 1:
 Are there actually any references to note 3 left, or can it be deleted?

SB comment:
 Tables 1-4 contain references to "note 3" which is a placeholder.

RC 1:
 SB comment:
 It looks like the gap has now been filled in by table 5, so change note 3 to link to table 5.

RC 1: Accepted, The Footnote was still needed to flag some numbers that had not arrived for v4.1. Bob M. and Justin M. are trying to help get the numbers from the IETF. The numbers needed are for NAS-Port and Service Type. The Footnote will be removed when these numbers are included.

SB resolution1/3/2003
 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	5.3.7.3
---------------	---------

Author: **Peter Ecclesine**

Comment Type: Editorial Vote: Disapprove Comment Status: Accepted Cmntr Response: Open

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
36	4	+02	RC1: Text says ", but should not be passed on to the old AP."	RC1: suggested_remedy = should read ", but should be passed on to the old AP."	RC1: Comment Accepted: Suggested Remedy applied.

Clause	5.4
---------------	-----

Author: **Peter Ecclesine**

Comment Type: Technical Vote: Disapprove Comment Status: Open Cmntr Response: Open

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
-------------	-------------	-----------	----------------	-------------------------	-------------------

26 13 39

RC1:
"In the last ballot, I submitted a comment that expressed concern about the trust model for AP to AP communications. The comment was declined.

Document 02/758 presented by Bill Arbaugh actually demonstrates a model whereby AP to AP communications is achieved through an acceptable trust model. The proposal in 02/758 doesn't presume AP to AP trust, the communications between APs are authenticated to ensure such trust."

SB comment:
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.

RC1:
suggested_remedy = Incorporate the mechanisms described in 02/758

SB remedy:
The comment contains the required changes

RC1 response:

SB response:
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	6.6
---------------	-----

Author: **Peter Ecclesine**

Comment Type: Technical

Vote: Disapprove

Comment Status: Open

Cmntr Response: Open

Page Line ID Comment

Suggested Remedy

Resolution

31 22 37

RC1:
"In the last Sponsor Ballot, it was suggested that the architecture should be revised (and possibly RADIUS removed) to enable fast and secure roaming. The comment was declined with the comment that I had not suggested a viable alternative and a reference to the another comment on the same topic.

If reliance on RADIUS is not removed then the draft must demonstrate clearly how fast and secure roaming is achieved using RADIUS. Alternatively, document 02/758 has shown a fast and secure mechanism that does not rely on RADIUS to secure context transfer."

SB comment:
Remove reliance on RADIUS and/or redesign architecture so that fast and secure roaming is possible.

RC1:
suggested_remedy = Add text showing how fast roaming can be achieved using RADIUS or add text based on the mechanisms in 02/758

SB remedy:
Add the messages indicated in the comment

RC1:

SB resolution:
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	Annex A
---------------	----------------

Author: **Arnoud Zwemmer**

Comment Type: Technical Vote: Disapprove Comment Status: Declined Cmnr Response: Open

<i>Page Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
------------------	-----------	----------------	-------------------------	-------------------

53 0 #02 RC1:

I disagree with adding the new MIB definitions introduced in draft 4.1. The original MIB was fine, but I believe this new MIB falls entirely outside the scope of the Task Group's PAR. It is not related to the goal of TGf, a recommended practice to enable multi-vendor interoperability over the DS. Enforcing single station association, communicating roaming of stations via an IAPP, and flipping switch tables are good examples of recommended practices within the task group's PAR, and adding a MIB to configure and monitor IAPP operation is also well within scope.

Yet adding an entire new SNMP configuration MIB for generic 802.11 operation is something completely different: this is like adding new network management functionality to Access Points for configuration and monitoring by SNMP network management stations. In fact, the MIB adds all kinds of 802.11-specific configuration and monitoring elements the entire TGf draft does not talk about. Furthermore, the management information is almost all corresponding to the 802.11 wireless interface itself (between STA and AP), which the Task Group should really not touch. This is really a change (namely an extension) to the basic 802.11 MIB of the 802.11-1999 standard and it is a technical change, which I believe the Task Group is not allowed to do.

So, while I would encourage extending the currently existing 802.11 MIB with more information, I feel this should be done in a separate Task Group and not in TGf. At this moment already, MIB objects are being standardized in TGe and TGi that overlap with the seemingly random set of objects TGf added in the latest draft. An example is the unicast cipher suite that is selected for each station. This is already defined in the TGi MIB, where it belongs. I suspect that TGe and TGi will work to define MIB objects corresponding to QoS and security behaviour, respectively. Other task groups will add their respective objects. This is something TGf should not interfere with.

RC1:
suggested_remedy = Remove the new MIB objects added in draft 4.1.

RC1: Comment Declined. The MIB was added as a direct result of comments from previous letter ballots. The MIB variables were determined in agreements with other Task Groups. TGf believes that these MIB variables are needed for reasonable operation.

Clause

Annex B

Author: **Peter Ecclesine**

Comment Type: Editorial

Vote: Disapprove

Comment Status: Open

Cmnt Response: Open

Page Line ID Comment

Suggested Remedy

Resolution

55 0 #02

RC1:
"The majority of the MIB parameters listed in Annex B are completely out of scope for TGf. Many of the parameters listed should be defined in other, more appropriate task groups (i.e. TGe, TGi, TGk).

Some examples include:

* dot11AddrTableEntryEncryption - the encryption mechanism used by the station in an AP that allows mixed encryption.

* dot11AddrTableEntrySignalStrength - the signal strength of the last frame received from the station in dBm.

* dot11AddrTableEntryLinkQuality - indication of the quality of the signal as measured in the last frame received from the station."

RC1:
suggested_remedy = Remove Annex B

RC1: Comment Declined. The MIB was added as a direct result of comments from previous letter ballots. The MIB variables were determined in discussions with members of other Task Groups. TGf believes that these MIB variables are needed for reasonable operation.

Moved to accept comment resolution as stated above : Butch Anton, 2nd: Tom T. 4 yes 2 no 1 abstain

Clause	General
---------------	---------

Author: **Arnoud Zwemmer**

Comment Type: Technical

Vote: Disapprove

Comment Status: Open

Cmnt Response: Open

Page Line ID Comment

Suggested Remedy

Resolution

RC 1:
Regarding comment ID 75: the commenter exchanged ideas with Bill Arbaugh on this topic. It seems our ideas are aligned. The basis for declining this comment (namely that it requires changes to 802.11) by the task group is eliminated if it does not require a message from a STA to an AP that it intends to roam.

SB comment:
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.

RC 1:
suggested_remedy = Instead of having a message from STA to AP to indicate a roam (as in the original comment's suggested remedy), distribute context information beforehand to a graph of neighbouring APs, similar to proposed in presentation 11-02-758r1 from Bill Arbaugh, thus enabling 'forward roaming'

SB remedy:
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.

RC 1:
SB resolution:
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.

Regarding comment ID 74: the commenter thinks he is misunderstood by the task group. The commenter agrees that RADIUS extensions are common and he is pleased that the task group finally makes this clear now while declining the comment. It is just that the commenter could not unambiguously derive this from the draft, whether IAPP backends for RADIUS servers would be necessary or that a standard RADIUS server will suffice. It seems logical at points in the draft where a special Service-Type is used (IAPP-Register), but for other standard RADIUS types (Call-Check) the draft causes confusion with more people than just the commenter that this is no different than a standard RADIUS request, which would imply it possibly is meant to work with any standard RADIUS server.

The security issue is not an issue if indeed IAPP-backends are required. Changed nature of comment to Editorial.

SB comment:

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?

suggested_remedy = The clarification that was added that address resolution can only be performed after having registered with the RADIUS server already helps, because this is IAPP-specific. Just add another sentence somewhere in the overview that the entity communicates with an IAPP-aware RADIUS server, to make clear from the beginning that it requires RADIUS extensions.

SB remedy:

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.

Comment Accepted: in Clause 1.3 add "The RADIUS server must provide extensions for IAPP specific operations."
Moved: Bob O 2nd Justin McCann Vote:
Unanimous

SB Resolution:

Extensions to RADIUS servers are a common occurrence 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 potential for a security issue mentioned; the access is not via MAC address only, but via MAC address and shared secret.

0 0 67 RC 1:
Sadly I don't expect us to ever agree on this one. I don't accept your resolution.

SB comment:
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.

RC 1:
SB comment:
This document should be put "on-hold" until a use that is identifiable to an end-user or network administrator is identified.

RC 1:
Comment declined same reasoning as SB --
The reason being that the comment is non-responsive per the Sponsor Ballot rules.

Sb comment:
Declined - the reason being that the comment is non-responsive per the ballot rules.

Author: **William Arbaugh**

Comment Type: Technical Vote: ADVISORY O Comment Status: Open Cmnr Response: Open

<i>Page</i>	<i>Line</i>	<i>ID</i>	<i>Comment</i>	<i>Suggested Remedy</i>	<i>Resolution</i>
0	0	97	RC1: The current IAPP protocol is reactive rather than proactive increasing the delay on REASSOCIATION by an order of magnitude. As requested by the TG, I will provide the full text proposal integrated into the current draft via email to the chair.	RC1: suggested_remedy = The protocol should be made proactive, perhaps optionally, to reduce the delay in support of fast roaming.	RC1: