ATNP/WG2/7
WP221
11 January 1996

AERONAUTICAL TELECOMMUNICATIONS NETWORK PANEL

WORKING GROUP TWO

Brisbane 5/2/96 - 9/2/96

Defects Found in Internet SARPs v3.1

and proposed solutions

Presented by Jean-Michel Crenais

Prepared by Stéphane Tamalet

SUMMARY

This paper presents questions and defects raised during a review of Internet Draft SARPs Version 3.1 , and proposes modifications.

DOCUMENT CONTROL LOG

SECTION	DATE	REV. NO.	REASON FOR CHANGE OR REFERENCE TO CHANGE
	11-Jan-96	Issue 1.0	

TABLE OF CONTENTS

1.Introduction

Validation of the draft ATN SARPs is continuing through simulations, ATN software implementation and implementation testing. In the scope of the EURATN software upgrade to conform to the CNS/ATM-1 Package SARPs, the review of the Draft SARPs version 3.1 raised a number of questions and defects. This paper presents a summary of the found defects, and recommends changes necessary to the draft SARPs Version 3.1.

2. « APRL for Air/Ground Route Initiation » section 3.5.2.13

2.1.Optional Support of optional non-use of IDRP

2.1.1.Problem Summary

The Draft SARPs section 2.4.3 mandates the support by **A/G routers** of the mechanisms necessary to support the optional non-use of IDRP: *« 2.4.3.b. ATN Routers of class 5 (Air/Ground Routers) and of class 7 (Airborne Routers without IDRP) shall also implement the mechanisms necessary to support the optional non-use of ISO 10747 »*

Remark: the end of paragraph 2.4.3.b is incoherent: « where the as specified in chapter 3 »should be replaced by « as specified in chapter 3 »

However, in the *general* (i.e. related to both airborne and air/ground routers) APRLs for Air/Ground Route Initiation (section 3.5.2.13.1), the support of the optional non-use of IDRP, covered by the **noIDRP** item, appears to be set optional:

noIDRP	Support of optional non-use of IDRP	3.5.2.11	O

Since the section 3.5.2.11 referenced by this item describes the Optional non-use of IDRP procedures for Airborne routers **and** Air/Ground routers, the noIDRP item is, with regard to the Air/Ground routers, in contradiction with Draft SARPs section 2.4.3.

2.1.2.Problem Resolution

It is believed appropriate to remove this inconsistency by change to the APRLs.

The following APRLs are proposed in replacement of the noIDRP APRL:

noIDRP-a	Support of optional non-use of IDRP (Airborne BIS)	3.5.2.11.2	O
noIDRP-ag	Support of optional non-use of IDRP (Air/Ground BIS)	3.5.2.11.1	M

In addition, all predicates referring to the *noIDRP* item in the next APRLs tables (section 3.5.2.13.2 to 3.5.2.13.5) should be replaced by a predicate referring accordingly either to *noIDRP-a* or *noIDRP-ag.*

2.2.Reversal of initiator and responder roles for the BIS-BIS connection compared with the Subnetwork connections

2.2.1.Problem Summary

For the purpose of ensuring the fastest route initiation procedure, Draft SARPs section 3.5.2.9.b requires that the role (initiator or responder) of a router in the establishment of a BIS-BIS connection be the opposite of its role in the establishment of the subnetwork connection. It is then required that: *« If the ISH PDU was received from a subnetwork connection which was established with the local ATN Router in the responder role, then the BIS-BIS connection shall be established with the* ***ListenForOpen*** *MO attribute set to false. Otherwise, the* ***ListenForOpen*** *MO attribute shall be set to true »*

The related items in the APRLs for Air/Ground Route Initiation in section 3.5.2.13, seem to be inconsistent with the procedure. Indeed, the description of the **initIDRP-*xx*** profile requirements gives to believe that the responder of the subnetwork connection must listen for OPEN BISPDUs and must therefore be the responder of the IDRP connection, and on the other end that the initiator of the subnetwork must not listen OPEN BISPDUs and so must initiate the IDRP connection.

The four questioned APRL items are the following:

In table of section *3.5.2.13.2 Airborne Router - Connection Responder:*

initIDRP-ar	IDRP startup procedures - ListenForOpen set to **true**	3.5.2.9	^noIDRP:M

In table of section *3.5.2.13.3 Airborne Router - Connection Initiator:*

initIDRP-ai	IDRP startup procedures - listenForOpen set to **false**	3.5.2.9	^noIDRP:M

In table of section *3.5.2.13.4 Air/Ground Router - Connection Responder:*

initIDRP-agr	IDRP startup procedures - ListenForOpen set to **true**	3.5.2.9	M

In table of section *3.5.2.13.5 Air/Ground Router - Connection Initiator:*

initIDRP-agi	IDRP startup procedures - listenForOpen set to **false**	3.5.2.9	^noIDRP:M

2.2.2.Problem Resolution

It is believed appropriate to remove this inconsistency by changes to the APRLs.

The following changes to the initIDRP-*xx* APRLs are proposed:

In table of section *3.5.2.13.2 Airborne Router - Connection Responder:*

initIDRP-ar	IDRP startup procedures - ListenForOpen set to ~~**true**~~ **false**	3.5.2.9	^noIDRP:M

In table of section *3.5.2.13.3 Airborne Router - Connection Initiator:*

initIDRP-ai	IDRP startup procedures - listenForOpen set to ~~**false**~~ **true**	3.5.2.9	^noIDRP:M

In table of section *3.5.2.13.4 Air/Ground Router - Connection Responder:*

initIDRP-agr	IDRP startup procedures - ListenForOpen set to ~~**true**~~ **false**	3.5.2.9	M

In table of section *3.5.2.13.5 Air/Ground Router - Connection Initiator:*

initIDRP-agi	IDRP startup procedures - listenForOpen set to ~~**false**~~ **true**	3.5.2.9	^noIDRP:M

3.Back-off procedure and ISO 8208 diagnostic codes

3.1.Problem Summary

The current text in draft SARPs section 3.5.2.2.1.1 mandates the use of the back-off procedure when a call is cleared with diagnostic code 133, 160..163, or 240,249. Additionally the note 4 in section 3.5.2.4, allows to know that when a call is cleared due to a call collision (i.e. code 130) the call collision resolution procedure must be applied by the SNDCF.

The text in SARPs can be illustrated by the following table:

Diagnostic codes			procedure to be followed
Diagnostic	Hex value	Decimal value	
	0000 0000	0	?
version number not supported	1000 0000	128	?
length field invalid	1000 0001	129	?
call collision resolution	1000 0010	130	Call collision resolution procedure
proposed Directory Size too large	1000 0011	131	?
LREF cancellation not supported	1000 0100	132	?

received DTE refused or received NET refused	1000 0101	133	back-off
invalid SNCR field	1000 0110	134	?
ACA compression not supported	1000 0111	135	?
V42bis compression not supported	1000 1111	143	?
inactivity timer expiration	1001 0000	144	Not Applicable
Need to re-use the circuit	1001 0001	145	?
by local means	1001 0010	146	?
invalid NSEL in received NET	1001 0011	147	?
	1010 0000 to 1010 0011	160 to 163	back-off
system lack of resources	1111 0000	240	back-off
	1111 0001 to 1111 1000	241 to 248	back-off
Connection Rejection - unrecognised protocol identifier	1111 1001	249	back-off

The discussions that took place, since Banff, between CCB members regarding ISO 8208 call clearing and the back-off procedure resulted in the identification of two questions which are considered to be incompletely answered by 95110019.CP, i.e. new section 3.5.2.2.1.1 :

1. For which diagnostic code should back-off be required ?
2. What action should be taken in response to call cleared with diagnostic codes not resulting in the use of the back-off ?.

Discussions on the first question among CCB members resulted in an informal proposal exchanged via the atn-technical list (see Annex A). It is proposed that WG2 decides on this question based on Annex A.

As regards the second point, recommendations should be added to the SARPs, concerning the procedures to be followed for all other codes (such as for code 130, for which use of the call collision resolution procedure is required (cf. note 4 in section 3.5.2.4). It is feared, otherwise, that omitting the specification of the appropriate procedures may lead to have implementation adopting an invalid procedure for certain of the clearing diagnostic code, such as retrying the call immediately whereas there is a high probability that the call will be cleared with the same diagnostic.

3.2.Problem Resolution

For the purpose of avoiding that an SNDCF overload the network with call requests which will never be accepted, it is proposed to add the following paragraph in the SARPs, just after the paragraph 3.5.2.2.1.1.a:

If the call is cleared with a diagnostic code reporting an error that the SNDCF is unable to correct and which would be anyway returned again if the call was retried, then the called DTE shall be removed from the polled DTEs list. Otherwise, the SNDCF shall retry the call after having modified the call user data and removed the cause of the previous call rejection.

4.Generation of the Leave Event on expiration of the inactivity timer

4.1.Problem Summary

A discussion between CCB members raised a question regarding the usefulness of having the leave event generated when the SNDCF receives a clear which indicates an inactivity timer expiration.

Those in favour of having the leave event generated on expiration of the inactivity timer argue that the inactivity timer expiration is after all an error condition and only occurs when communication has effectively been lost with the remote DTE, either because of problems in the network, or perhaps because the aircraft went out of range without logging off and generating a real leave Event.

At first glance, it seems correct to assume that the inactivity timer can never elapse whereas Air/Ground communication remains possible, since either Keepalive BISPDUs or ISHs PDUs are assumed to be periodically exchanged and thus prevent inactivity timer expiration.

However, certain cases have been found, for which it seems possible to have the inactivity timer expiring due really to inactivity and not to error condition.

In the case, for instance, of an IDRP-equipped airborne router connected with a unique A/G BIS via several mobile subnetworks (e.g. SN1, SN2 and SN3); it is assumed that the IDRP connection is established and keepalive BISPDUs are exchanged periodically. These keepalive BISPDUs will always be transmitted via the same mobile subnetwork (e.g. SN1). Then, since IDRP is in established state, the periodic transmission of ISH PDUs has been suppressed as recommended in Draft SARPs section 3.5.2.9; and, as a result, in the absence of application traffic going through SN2 or SN3, the inactivity timer for SN2 or SN3 connections may elapse whereas communication may be possible.

On the other hand, there are cases where the inactivity timer will not elapse although there are problems in the network preventing data units to be exchanged: Indeed, this is because this timer is reset every time a data unit is transmitted on the VC and not when data units are received on the VC. As a consequence, the inactivity timer does not elapse if data units are periodically sent on the VC, even when the communication is effectively lost and no data unit is received at the other end.

It is therefore believed that the inactivity timer cannot be used with enough confidence to detect loss of communication and that it is only useful to clear unused VCs.

Anyway, if the inactivity timer is used with mobile subnetwork connection, the following should be considered: if a leave event is generated on expiration of the inactivity timer, the related configuration information (the ISH) of the neighbour BIS on this subnetwork is removed from the FIB and consequently, a routing initiation procedure will have to be re-started for having a new VC opened between the 2 BISs. Then:

- If the subnetwork is a mobile subnetwork which does provide Connectivity Information, a VC will be immediately re-established on receipt of the Join event. The use of the inactivity timer for dropping unused VCs, serves nothing since the unused VCs are immediately re-established. On such type of mobile subnetwork it is therefore believed that either the inactivity timer must not be used or if used that no leave event must be generated when the VC is cleared due to this timer.

- If the subnetwork is a mobile subnetwork which does not provide Connectivity Information, the DTE polling mechanism is used. Then, if the DTE for which the inactivity timer elapsed is not subject to a back-off period, then the VC is immediately re-established and the use of the inactivity timer for dropping unused VCs, serves nothing. On the other hand, if it is subject to a back-off period, then the use of inactivity timer appears to be unacceptable since, due to it, communication can be lost for the duration of the backoff period . On such type of mobile subnetwork it is therefore also believed that either the inactivity timer must not be used or if used that no leave event must be generated when the VC is cleared due to this timer.

Considering the case where the use of the inactivity timer is permitted with the recommendation to not generate any leave event on expiration of this timer, the following scenario could be observed: the inactivity timer expires, whereas the communication remained possible; the VC is cleared but no leave event is generated; the aircraft comes then out of range without logging off and there is an effective loss of communication; In such a case, the routers will not be warned immediately that communication via the subnetwork is no more possible; they will only detect the effective loss of communication when an attempt to open a new VC over the subnetwork is made. There is a risk of long lived routing black hole: the A/G BIS will advertise an invalid route until someone attempts to follow this route. As soon as an NPDU will attempt to follow the route, the A/G (or airborne) Router will not succeed in opening the subnetwork connection and will therefore be warned.

Since the loss of connectivity is detected on the first NPDU attempting to follow the invalid route, it can be considered that the routing black hole has a limited consequence: only the first NPDUs are lost. But whether or not this is acceptable, this is another question which is proposed to be discussed by the group.

4.2.Problem Resolution

The following two solutions should be considered:

- either the use of inactivity timer is precluded with mobile subnetwork

- or it is permitted with the implementation of a specific procedure in the mobile SNDCF, which will not generate the leave event when the VC is cleared due to inactivity but which on the other hand will generate a leave event if a subnetwork connection previously closed due to inactivity fails to be re-established.

WG2 should consider which of the two solutions is preferred.

5.Comment on the recommendation to suppress ISH periodic transmission when a BIS-BIS connection has been established

5.1.Problem Summary

Draft SARPs section 3.5.2.9 recommends to suppress the periodic transmission of ISH PDUs when a BIS-BIS connection has been established, except when a watchdog timer is applied to the subnetwork connection.

Since this is only recommended and not required, the following situation can be encountered: an IDRP-equipped Airborne router implementing this recommendation is connected with an Air/Ground Router not implementing the recommendation. Then, as soon as the IDRP connection is established between these routers, the Airborne router stops transmitting ISHs while the Air/Ground Router periodically transmits its ISH.

It is believed that there is a risk that the Air/Ground router chooses to remove the airborne ISH from its local FIB at expiration of the ISH Holding Time (and consequently closes the IDRP connection).

5.2.Problem Resolution

It should be made clear that when a BIS-BIS connection has been established, the configuration information (i.e. the ISH) of the adjacent router must be considered valid for the duration of the IDRP connection.

A first solution could be to require that the Information Holding Timer be stopped when the associated BIS-BIS connection is established and not restarted until the IDRP connection is closed. This solution presents however the following two drawbacks:

1. it necessitates the implementation of an ATN specific procedure controlling the stop and restart of ES-IS Holding Timers

2. it is difficult to know the appropriate value to be used for the Holding Timer when this timer is restarted.

Another solution could be to require that a router stopping ISH retransmission must first transmit an ISH with a very long Holding Time, an Holding Time which exceeds the duration of the IDRP connection. Considering that the two bytes HT field in ISH PDUs allows to convey an Holding Timer value of up to 65535 seconds (i.e. more than 18 hours), and which therefore exceeds the duration of mobile connections, this solution seems reasonable. It presents however the following two drawbacks:

1. it necessitates the implementation of an ATN specific procedure performing the transmission of an ISH with a very long Holding to the peer BIS with which the BIS-BIS connection has been established

2. it requires that the ES-IS implementations support large range of values for the Holding Time field in transmitted PDUs (cf. Htv item of ISO 9542 APRL in section 8.2.2)

This second solution is preferred considering that it allows the routers to respect the choice of procedure made by the peer routers: a watchdog timer will monitor connections established with peer routers which do not suppress the periodic transmission of ISHs; and no watchdog timer will run for connection with peer routers which suppress the periodic transmission of ISHs.

It is proposed to add the following text at the end of section 3.5.2.9:

> *3.5.2.9.e. If the periodic transmission of ISH PDUs is suppressed by the local BIS when a BIS-BIS connection has been established, the IS-SME shall issue to the peer BIS and as soon as the BIS-BIS connection is established, an ISH PDU the Holding Time of which is set to 65534 seconds.*

Note: 65535 seconds is not proposed as value for the Holding Time in this case, because certain implementations use this specific value (0xffff in hexadecimal) to configure infinite Holding Times.

6.ATN Use of Priority

6.1.Problem Summary

Compared to section 2.6, ATN Priority Provisions, of draft SARPs Version 3.0, the new text in draft SARPs Version 3.1 on the ATN Use of Priority (section 2.8) differs, amongst other things, by having the Application priority mainly mapped to the Transport Priority.

It was understood with the previous text (draft SARPs 3.0) that the application priority had first to be mapped to the network priority and possibly, if used, was then mapped to the transport priority.

The way section 2.8 is drafted (section 2.8.1.b) implies that the support of Transport priority is mandatory. This is inconsistent with transport APRL, section 5.2.4,1.1.2, which leaves the support of the Transport priority as optional (ATN14 and ATN22 items).

6.2.Problem Resolution

Since the use of the Transport priority parameter is straightforward, and since it seems to be the easiest way to map Application priority to Network priority, it is proposed to mandate the transport priority by making the following changes in chapter 5:

in section 5.1.2 d), replace " the application Service Priority to be mapped into the resulting CLNP NPDUs according to table 2.2; " by " the application Service Priority to be mapped into the resulting transport priority according to Table 2.3; "

in section 5.2.4.1.1.2, items ATN14 and ATN22 : in the support column, replace O by M

7.Questions regarding the ATSC Class

7.1.Problem Summary

The ATSC Class is used to convey the strong Quality of Service requirements. Currently, only the Transit Delay semantics of ATSC Class are defined in the Draft SARPS.

Currently, the only semantic hidden behind the term ATSC class is the Transit Delay, with the maximum end to end transit delay values proposed in table 2.2 .

But it is not clear how IDRP will manage this attribute except if an assumption is made that transit delays on ground subnetworks are negligible.

Indeed, if it is not assumed that ground to ground transit delays are negligible, IDRP on ground routers will have, when advertising a route received from an adjacent BIS to another adjacent BIS, to update the ATSC class of the route, based on:

- the knowledge of the maximum transit delay experienced between the local BIS and the adjacent BIS to which the route is advertised, and
- the previous ATSC class of the route.

As an example, a ground BIS may have to re-advertise a route received from an adjacent BIS to another adjacent BIS with which the experienced BIS to BIS maximum transit delay is 2 seconds. If the received route included an ATSC Class security tag of value C (meaning maximum end to end transit delay is not greater than 13 seconds), the local router should re-advertise a route which ATSC class corresponds to " Class C + 2 seconds ".

The question then is: Is " Class C + 2 seconds " equal to Class C or to Class D ?

It is believed that the way the ATSC class is encoded at IDRP level will not allow IDRP to answer this question; IDRP will then choose randomly to downgrade or not the previous ATSC class of the route. This leads to have strong routing requirements based on very approximate QoS estimates.

7.2.Problem Resolution

It is proposed that the value of the maximum end to end transit delay be encoded in the ATSC Class Security Tag Set of IDRP UPDATE BISPDUs, instead of the single letter 'A' through 'H' identifying the ATSC class.

It is therefore proposed to replace the current text in section 8.3.1.3.2 by the following text:

> *The Tag Set Name shall be set to 0000 0110 and the Security Tag is always one octet in length. When this security tag set is present, the Security tag shall identify the maximum remaining transit delay that would be experienced by SNSDU size of 512 octets to reach the NLRI advertised in the UPDATE PDU. The value encoded in the Security Tag shall be a positive integer. The maximum transit delay is specified in units of 1 sec.*

The current text in section 8.3.1.4.2 is proposed to be replaced by the following text:

> *When a route is advertised to an adjacent BIS, and according to local policy rules specified by a System Administrator and dependent on the BIS to which the route is being advertised and the route's NLRI, then:*
>
> > *a) if a BIS advertises a route whose destinations are located in its own Routing Domain, then the originating BIS may add an ATSC Class security tag to the route identifying the maximum transit delay between the advertised NLRI and the BIS to which the route is being advertised*
> >
> > *b) When a BIS re-distributes a route which has been learned in an UPDATE PDU that contains the ATSC Class security tag, it shall update the value of the Security Tag before advertising the route to another BIS. The updated value shall be computed by adding the maximum transit delay experienced between the local BIS and the BIS to which the route is advertised, to the value of the ATSC class security tag that was received in the UPDATE PDU*

The paragraph c) in section 8.3.1.6.3 is proposed to be replaced by:

> > *c) When an ATSC Class security tag occurs in all component routes then the aggregated route shall contain an ATSC Class security tag. The maximum transit delay encoded in the ATSC Class security tag of the aggregated route shall be the longest maximum transit delay encoded in the ATSC Class security tag of the aggregated route's component route. If an ATSC Class security tag is not present in any component route then the aggregated route shall not contain an ATSC Class security tag.*

8. Security classification

8.1.Problem Summary

In Draft SARPs 3.1 section 2.7, note 4 explains that the ATN security mechanisms " *uses a Security Label in the header of each CLNP Data PDU to convey information identifying the traffic type of the data and the application's routing policy and/or strong QoS Requirements. No mechanisms are provided to protect the integrity of this label or its binding to the application data ".* In section 2.7.1, it is then stated that " *An NPDU's Security Label shall identify the traffic type of its user data "* (2.7.1.c) and that this traffic type shall be expressed through further information encoded into the security label, as, an ATSC class, or an ordered preference for the types of air-ground subnetwork to be used, or no routing policy preference

On the other hand, in chapter 6 (section 6.2.2.5), the ATN Security information is described as being made of 2 possible Tag Sets: the Traffic Type and the security classification.

Although it is then explained that the purpose of the security classification field in the ATN security label is to permit a later extension of the ATN to handle classified data, it is found inconsistent that no reference to the Security Classification be made in chapter 2, since it is described in section 6.2.2.5.8.

Since the security classification field is defined " for future use ", its definition in the Package-1 SARPs is questionable.

8.2.Problem Resolution

It is proposed to remove the inconsistency existing on this subject between chapters 2 and 6, by removing the definition of the security classification tag set from chapters 6 and 8.

ANNEX A

Editor's note :

The following material is based on an informal proposal made via e-mail on the atn-internet-technical *mailing list in December 95. It has been inserted as an annex to WG2/7/WP221 in order to help and possibly accelerate the decision-making on the issue of the ISO 8208 diagnostic codes and the use of the « back-off » procedure.*

The author of this proposal was J.M. Crenais.

This paper explicitely name various WG2 experts who participated to the discussions which occurred about the technical issue raised in the proposal. All references to opinions or statements made by these experts have been included without their previous approval and should therefore be considered as the pure intrepretation of their thoughts from the author.

May these experts forgive the author who only dared take this decision for the sake of our SARPs !

Proposal to solve the issue about the use of the « back-off » procedure and the ISO8208 diagnostic codes.

Since Banff, several messages and CCB VRCIs were issued about the use of the « back-off » procedure in the airborne router, depending on the value of the ISO 8208 diagnostic code contained in the Clear Indication packet.

The following people and documents were involved in the discussion:

1. Based on Action 6/10 from Banff, Tony Whyman submitted 95110060.DR. In summary, the DR was requesting the identification of call clearing diagnostic codes that result in the suppression of the « back-off » procedure.

2. In response to DR60, Tony submitted 95110019.CP, based on Flimsy #10 from Banff,

3. On 31/10/95, Ron Cossa sent an e-mail message containing a Word 6.0 document which contains comments on Flimsy #10,

4. On 7/11/95, Helen Thulin sent an e-mail message containing new proposals on this issue.

5. A first version of this Annex to WP221 was released, but it generated no further comments and 95110019.CP was finally accepted and implemented in version3.1 of the SARPs as new section 3.5.2.2.1.1.

However, it is still believed that the text proposed in 95110019.CP does not solve the problem. The purpose of the present document is therefore to clarify the discussion on this issue and to make proposals so that WG2 can correct definitively the SARPs on this issue.

The following table lists Helen's and Tony's proposals (as understood by the author !) as regards the use of the « back-off » procedure based on the value of the diagnostic codes.

Note 1.— In CP19 (i.e. section 3.5.2.2.1.1), only the values for which the « back-off » procedure is required are listed; nothing is said as regards the other values. I have interpreted this as follows: the « back-off » procedure shall NOT be implemented when the diagnostic codes have any other value. This is reflected in the following table.

Note 2.— Another Change Proposal (ref. 95110018.CP) suggested to split code 133 [1000 0101] in 2 codes:

- *code 133 [1000 0101], meaning « Received DTE refused, or received NET refused »,*
- *code 147 [1001 0011], meaning « Invalid NSEL in received NET ».*

Since Helen endorsed this proposal in her e-mail (despite a typo: 1000 0011 instead of 1001 0011), and nobody else disagreed (Klaus-Peter rejected 95110018.CP for other reasons), I have assumed this change was ACCEPTED and it is reflected in the following table.

Diagnostic codes			Use of the « back-off » procedure	
N° in Table 7-6	Hexa value	Decimal value	as proposed by Tony	as proposed by Helen
12	0000 0000	0	NO	YES
2	1000 0000	128	NO	NO
3	1000 0001	129	NO	NO
4	1000 0010	130	NO	NO
5	1000 0011	131	NO	NO
6	1000 0100	132	NO	NO

7	1000 0101	133	YES	YES
8	1000 0110	134	NO	NO
9	1000 0111	135	NO	NO
10	1000 1111	143	NO	NO
13	1001 0000	144	NO	NO
14	1001 0001	145	NO	YES
15	1001 0010	146	NO	YES
16 *(see Note 2)*	1001 0011	147	NO	YES
not listed	1010 0000 to 1010 0011	160 to 163	YES	*Nothing mentionned*
11	1111 0000	240	YES	YES
not listed	1111 0001 to 1111 1000	241 to 248	YES	*Nothing Mentionned*
1	1111 1001	249	YES	NO

A quick check shows that Helen and Tony disagree on 5 codes: 0, 145, 146, 147, and 249.

Note.— I consider that Helen would agree with Tony on codes 160 .. 163, and 241 .. 248 (In France, we have a proverb which says: « Who says nothing, means Yes !»...).

At first, according to Helen, the « back-off » procedure should only be performed when « (1) the ground DTE refuses the connection due to a saturation of resources or when (2) the airborne DTE/NET is not accepted by the ground DTE due to local policies ». I agree with her first point, and consequently I support to run the « back-off » procedure in response to Clear Indications containing codes **0, 133, 145, 146, and 240** as she proposes in her e-mail. But I am questioning the second point which seems to be against the principle of « maximum connectivity » on the air/ground segment, and consequently I would prefer to NOT run the « back-off » procedure for code **147**.

Secondly, still according to Helen, the « back-off » procedure should not be started when « the call is rejected due to a subnetwork or ground SNDCF reported invalid format/content ». In principle, I agree with Helen, but in some of these cases, she also proposes to stop the polling mechanism (codes 128, 129, 132, 134, 135, 143, 249), and I am not sure about this: isn't it possible to have some implementations which could correct this invalid format/content, and consequently send an acceptable Call Request on the next try ?

Consequently, I would support to NOT run the « back-off » procedure for codes **128 ..132, 134, 135, 143, 144, and 249**, but to NOT say anything about stopping the polling mechanism.

As regards codes 160 to 163 and 241 to 248, I agree with Tony: the « back-off » procedure shall be run.

This would lead to the following table:

Diagnostic codes			Use of the « back-off » procedure as proposed by ...		
N° in Table 7-6	Hexa value	Decimal value	Tony	Helen	Jean-Michel
12	0000 0000	0	NO	YES	YES
2	1000 0000	128	NO	NO	NO
3	1000 0001	129	NO	NO	NO
4	1000 0010	130	NO	NO	NO
5	1000 0011	131	NO	NO	NO
6	1000 0100	132	NO	NO	NO
7	1000 0101	133	YES	YES	YES
8	1000 0110	134	NO	NO	NO
9	1000 0111	135	NO	NO	NO
10	1000 1111	143	NO	NO	NO

13	1001 0000	144	NO	NO	NO
14	1001 0001	145	NO	YES	YES
15	1001 0010	146	NO	YES	YES
16 *(see Note 2)*	1001 0011	147	NO	YES	NO

not listed	1010 0000 to 1010 0011	160 to 163	YES	*Nothing mentionned*	YES
11	1111 0000	240	YES	YES	YES
not listed	1111 0001 to 1111 1000	241 to 248	YES	*Nothing mentionned*	YES
1	1111 1001	249	YES	NO	NO

Helen also raises the fact that the « back-off » procedure should not be run in case of emergency calls. I agree, but I am wondering how an emergency call is actually distinguished from an « ordinary » call ? Is it via the priority parameter (CLNP values '12' and '13' from Table 2-2) ? Some text should probably be added to clarify this in the SARPs ... (Sorry ! I have no concrete proposal for this.)

In conclusion, I would propose to reject 95110019.CP and accept the following Change Proposal:

1. Replace the first paragraph of section 3.5.2.2.1.1 with the following text:

Whenever a Clear Indication is received in response to a Call Request that indicates rejection by the called DTE and includes a call clearing diagnostic code of 0, 133, 145, 146, 160 . . 163, or 240 . . 248, then the Airborne Router shall implement a « back-off » procedure. The « back-off » procedure shall comprise the effective quarantining of the called subnetwork address for a period configurable on a per subnetwork basis from 5 minutes to 20 minutes. During this period, a Call Request shall not be issued to the subnetwork address.

The « back-off » procedure shall not be started on receipt of a Clear Indication which includes any other call clearing diagnostic code.

Note 1.— Certain call clearing diagnostic codes in the range 128 . . 143 are used by the mobile SNDCF specified in Chapter 7. The semantics of these codes is described there in Table 7-6.

2. Original *Note.* in section 3.5.2.2.1.1 becomes: *Note 2.*

3. Replace Table 7-6 with the following table:

	Hexadecimal value	Decimal value	Diagnostic
1	1111 1001	249	Connection Rejection - unrecognized protocol identifier in user data
2	1000 0000	128	Version number not supported
3	1000 0001	129	Length field invalid
4	1000 0010	130	Call Collision Resolution
5	1000 0011	131	Proposed Directory Size too large
6	1000 0100	132	Local Reference Cancellation Not Supported
7	1000 0101	133	Received DTE refused, or received NET refused
8	1000 0110	134	Invalid SNCR field
9	1000 0111	135	ACA compression not supported
10	1000 1111	143	V42bis compression not supported
11	1111 0000	240	System lack of resources
12	0000 0000	0	Cleared by System Management
13	1001 0000	144	Idle Timer expiration
14	1001 0001	145	Need to re-use the circuit

15	1001 0010	146	By local means (to be used for system local error)
16	1001 0011	147	Invalid NSEL in received NET