ATNP/WG2
IP/497a

EUROCONTROL

Fault Management Requirements Analysis

Prepared by A. Whyman

Presented by M.G. Adnams

AERONAUTICAL TELECOMMUNICATIONS NETWORK PANEL

Working Group 2

Hawaii 18-22 January 1999

SUMMARY

This paper provides a proposed fault management concept. It is concerned with identifying the typical faults that might occur in an internet, how the source of the fault is diagnosed, and thereby deriving the requirements for Managed Objects and tools in support of Fault Management.

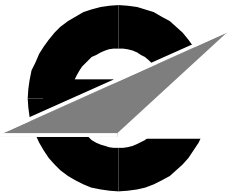

EUROCONTROL

ATN PROJECT

Fault Management Requirements Analysis

ATN Ref. : DED6/ATNCT/ProATN_Sup/DCI/AW_43

Author : Tony Whyman

Rev. No. : Draft 0.1

Date : 13/01/98

DOCUMENT CONTROL LOG

SECTION	DATE	REV. NO.	REASON FOR CHANGE OR REFERENCE TO CHANGE
	13/01/98	Draft 0.1	

TABLE OF CONTENTS

1. Introduction

1.1 Scope

This paper is concerned with identifying the typical faults that might occur in an internet, how the source of the fault is diagnosed, and thereby deriving the requirements for Managed Objects, Implementations and tools in support of Fault Management.

1.2 Purpose of Document

This document has been prepared as a discussion document for use within Eurocontrol and in projects implementing Systems Management services.

1.3 References

1.	ICAO DOC9705	ATN SARPs
2.	DED6/ATNCT/SYSMAN/ DCI/L7-13V1_2.doc	ATN Fault Management Requirements
3.	Version 1.4	ACI and ProATN: Convergent MIB
4.	ATNP/WG2/WP-478	Preliminary Draft Version 1.0 ATN Systems Management – Concept of Operations

2. Summary

2.1 Approach to the Analysis

In order to determine the requirements for Systems Management Managed Objects and Tools, it is first necessary to understand what a Network Manager needs to do to recognise, diagnose and repair faults. Key to this is an understanding of the recognition process by which a Network Manager learns about the existence of a fault – real or potential. The diagnosis and repair can then follow a logical exposition based upon experience and known techniques. It is thus an understanding of the "triggers" – the events – that cause a Network Manager to enter the fault diagnosis and repair procedure that is the subject of the first part of the paper.

A "Black Box" approach is taken for this part of the analysis. That is, ATN components, such as Routers and Subnetworks are viewed as "Black Boxes" and only externally visible faults are considered – that is faults which result in a failure to meet the component's service requirement. The simple rationale for this, is that only when a fault is externally visible can it be a problem; a fault that is not externally visible may the harbinger of an externally visible fault, and properly the subject of a background preventative maintenance process. However, it is the analysis of the externally visible faults that determines the "triggers" for Network Manager action, and by tracing such faults back to the underlying reasons, it is also possible to determine the proper preventative maintenance procedures. This paper also includes the identification of possible preventative maintenance procedures and identifies "derived triggers" that can result from such procedures.

Earlier work **[2]** has concentrated on identification of the detailed internal faults. This has been used in order to validate the identification of the externally visible faults by relating each such detailed case to the externally visible faults. It is also provides the basis for a determination of the underlying fault conditions that should be monitored for fault prevention.

The ATN components considered are:

1. Subnetworks

2. Routers

3. The end-to-end Transport Service.

Both subnetworks and Routers are tangible components and are the components of the Internet. However, the end-to-end Transport Service is more abstract in nature. On the other hand, it represents the user view of the ATN Internet and it is, essentially, the ATN seen as a whole system. Faults, when they are externally visible to an end user, are clearly very significant and the above is really saying that the Network Manager will react when there are visible faults in the main internet components (Subnetworks and Routers) and when a fault is visible to an end user.

The service that a subnetwork provides to its user is simple to define. It is the transport, across that subnetwork, of some unit of user data (i.e. a message or a packet) to another subnetwork service user, in line with the declared Quality of Service (including transit delay, undetected error probability and availability). This may be performed in the context of a virtual circuit (connection mode) or as discrete datagrams (connectionless). The externally visible failure modes are the possible deviations from this service.

Similarly, an ATN Router provides a simple packet forwarding service, with a declared probability of forwarding a packet along the optimal path to its destination whilst maintaining

data integrity. The failure modes of an ATN Router are thus very simple and limited and, not only that, the ATN Internet is actually very tolerant of such faults, as an inappropriate forwarding decision by one router is often repaired by a router further down the path.

The end-to-end transport service also provides an end-to-end service with a similarly limited list of fault conditions. The connection mode transport service provides reliable sequenced data transfer in line with the declared Quality of Service, and the connectionless transport service provides a simpler datagram service with a limited, but good, probability of delivery, again defined by a declared Quality of Service[1]. The connection mode transport service should be able to mask most ATN Router and Subnetwork faults from the end user. Only when they are so severe that the Quality of Service cannot be maintained are they visible to the end user.

Also, implicit in the ATN design is that in order to maintain the operational service to a high level of availability, it is important that there exist:

- Alternative data links to replace failed data links
- Alternative ATN Routers to replace failed ATN Routers.

This is about procurement of data links and physical deployment. Capacity planning and network design models will need to be developed to determine the data links and routers needed to meet demand, and the excess capacity and redundancy in the network needed to meet the availability target, given the *Mean Time Between Failures* and the *Mean Time To Repair* of the data links and routers.

There is thus a dependency on Performance Management for the assurance of long term network service availability.

2.2 Application Errors

It should be noted that the ATN Applications are not themselves considered in this analysis except when they are required to report ATN Internet Failure Modes. The principle is followed that an ATN Application will report application errors to its service users rather than to a Network Manager on the assumption that it is the Application User that is responsible for application level fault recovery and not the Network Manager. The Network Manager is responsible for ATN Internet Fault Management only.

2.3 Certification

A certification issue should also be noted as resulting from this analysis. The ATN Internet is a fault tolerant environment and fault conditions in Routers and Subnetworks can be tolerated without affecting the end user service. It thus follows that confidence in their correct operation does not have to be excessive. The end user Quality of Service can be maintained to high levels even with a relatively low confidence in the proper operation of ATN Routers and Subnetworks as long as there is sufficient redundancy of components and that common mode failure conditions are avoided. As Routers operate autonomously and no two Routers can be expected to have identical Routing Information and State History, common mode errors are

1. ———————————

[1] With a connectionless network, there is no concept of a *requested* Quality of Service. A user of the service gets the available Quality of Service and a Service Provider is expected to *declare* the Quality of Service that a user should reasonably expect.

unlikely to affect well tested Routers. With Subnetworks, physical level analysis will be needed to avoid sharing of physical resources (e.g. cables) between otherwise independent services.

On the other hand, there is a strong dependency on the correct operation of the end-to-end transport protocol for this happy situation to be maintained. Maintenance of the end-to-end Quality of Service (i.e. the avoidance of faults in the end-to-end service) is dependent on a well tested, well designed and high quality implementation of the transport protocol.

2.4 Identification of Fault Management Requirements

From an analysis of the possible “Black Box” failure modes, the triggers can be identified. This paper has identified these trigger events and gone on to discuss the procedures that a Network Manager will follow in order to identify the actual source of the fault. Often this will be in a different system to that reporting the fault. For example, most packet forwarding errors that are not correctable by another router will result in the eventual discard of the packet once it exceeds its lifetime count. The router that performs the discard may be many hops away from the one that mis-routed the packet. Only by tracing back along the route and looking for common path segments between many such discards can a Network Manager determine the router that is causing the problem. This analysis has thus concluded (unsurprisingly) that the CLNP Echo PDU is a primary tool in fault determination.

The diagnosis procedure may often require co-operation between the Network Managers of multiple domains. This has already been discussed in **[4]** and is not considered in detail here.

Once candidates for the source of the fault have been identified, there will need to be inspection of operational information (e.g. Routing and Forwarding Information) and statistics (e.g. data transfer counts) in order to provide diagnostic determination of the source of the fault. Fault repair may require short term isolation of the failed component. In the case of a router software fault, reboot provides a short term repair, with post-mortem analysis of the router state being required for correction of the “bug”. Other failure modes may require the intervention of an engineer. The need to isolate components before and during repair places a strong emphasis on the need for redundancy of such components if the end-to-end service is to be maintained during such outages.

This paper concludes by identifying first the requirements for Managed Objects and Network Management tools, as derived from this analysis, and then the formal identification of the Managed Objects, referring to the current working document on Managed Object Requirements **[3]**.

3. Failure Mode Analysis

From the point of view of fault analysis, the ATN is seen as comprising three main components:

- Subnetworks.
- ATN Routers
- The ATN Internet (providing the end-to-end transport service)

Each of these has failure modes which need to be considered, in turn.

3.1 Subnetworks

An ATN Subnetwork may provide either a virtual circuit type service or a connectionless service. In the context of the provided service, it is required to transport data units (packets) from one subnetwork service user (the source) to another unambiguously identified subnetwork service user (the destination), whilst meeting the required Quality of Service. The path between any two subnetwork service users is regarded as a “data link”. The failure modes of such data links can be readily identified as:

1. Total loss of the Communications Path (including complete subnetwork failure)

2. Partial loss of the Communications Path (loss of some packets but not all)

3. Corruption of packets, while in transit

4. Mis-delivery of packets

5. Packet re-ordering

6. In addition, the ATN includes mobile subnetworks and these introduce the additional failure mode of failure to join a subnetwork.

Note that not all network types will exhibit all possible failure modes.

3.1.1 Failure Mode #1: Total loss of the Communications Path

Loss of communications will, and indeed must, be detected by the Subnetwork Service User i.e. the attached End System or Router. In either case, this will be reported to the routing function, and immediate recovery should be by replacing the failed data link with an alternative path - if one exists. This may involve an exchange of routing information.

For repair of the failed data link, a Network Manager must be notified of the failure and a repair performed by the service provider. The notification of unexpected subnetwork connection loss by an End System or Router is thus a trigger event on which a Network Manager will start a fault diagnosis and repair process.

If the fault is transient, then recovery of the lost data link may be automatic or require manual intervention depending upon the local policies of the Network Manager. A permanent fault will only be cleared on repair.

3.1.2 Failure Mode #2: Partial Loss of Communications

Partial loss of communications cannot, in general, be detected by the ATN Routers; it is only detected by the end-to-end transport protocol recognising packet loss. On a data link, such as frame relay, packet delivery is not guaranteed (only a minimum bit rate is actually guaranteed) and hence loss of some packets cannot be regarded as an error. However, there does need to be some comparison of the send and received data rates at each end of (e.g.) a Frame Relay data link to ensure both contractual compliance by the service provider, and to possibly predict future loss of service. The error rate counts may also prove useful.

With a reliable network service, such as X.25, packet loss should be detected and recovered from by the subnetwork itself. However, if such losses were undetected, the situation is as above. Again, it will be important to compare sent and received packet counts to check for undetected packet loss by the X.25 Service Provider.

However, it should be noted that the a given subnetwork connection may between two Management Domains. In this case, there will also need to be co-operation between Network Managers in order to share such data.

In conclusion, this failure mode does not itself provide any externally visible events. It is rather some sort of “stealth” error condition that can only be detected by a diagnostic procedure that compares send and receive packet counts at either end of a subnetwork connection.

3.1.3 Failure Mode #3: Corruption of Packets

All ATN subnetworks are expected to provide CRCs to detect corrupt packets. However, it is possible for some error patterns to remain undetected by the CRC. Also, the CRC is sometimes regenerated in the subnetwork and this process can also introduce undetected errors.

If such an error occurs within the CLNP header, then this should be detected by the CLNP Header checksum (although this depends upon the originating End System generating a non-zero checksum in the first place, which is itself optional). A failure to validate the CLNP header checksum will cause the packet to be discarded. If the checksum is not available, then a CLNP packet with an error in the header may continue to be routed. It is probable that the packet will then be mis-routed, as a result, and eventually discarded. From this it may be concluded that the CLNP checksum feature is very desirable as it will detect such an error at source and hence save time in fault diagnosis.

If such an error occurs in the user data, then this should be detected by the transport header checksum on delivery at the destination End System.

Over air/ground datalinks using Deflate based compression, such a fault may also be detected by the additional packet level checksum used for compressed PDUs.

If this fault is transient, then end-to-end retransmission should be sufficient to recover from the fault, albeit with a consequential increase in transit delay. If the fault is permanent, then this could result in a “routing black hole” until routing information exchanges are also lost, resulting in the data link being declared as failed and an alternative path chosen.

In this case, the trigger events for Network Manager investigation are:

1. CLNP packet discard reports,

2. SNDCF Deflate checksum error reports

3. Transport level checksum error reports

4. Excessive TPDU retransmissions.

It should be noted that transport checksum errors cannot be assigned to individual transport connections and hence retransmission counts (which can be available on a per transport connection basis) will be needed to help determine which transport connections are affected.

3.1.4 Failure Mode #4: Mis-delivery of packets

It is possible for a subnetwork to deliver a packet to the wrong destination. This failure mode can be difficult to distinguish from router mis-routing (see 3.2.3) but should always be considered when there is no error in the Routing Information Base (RIB). Comparison of the sent and received data counts on the subnetwork connections may indicate the actual problem.

This is another example of a “stealth” error condition that can only be detected by a diagnostic procedure that compares send and receive packet counts at either end of a subnetwork connection.

3.1.5 Failure Mode #5: Packet re-ordering

A subnetwork may deliver packets in a different order to that in which they were sent. In some network types, this is a feature and not an error. In the ATN, it is the responsibility of the end-to-end transport protocol to detect out of order packets and to re-order them into their correct sequence. Hence, this should not result in a user visible fault.

However, if the network is X.25 or similar and performs segmentation on CLNP packets, then a mis-ordering could corrupt the re-assembled packet. The result will then be the same as packet corruption, as discussed in 3.1.3.

3.1.6 Failure Mode #6: Failure to Join a Mobile Subnetwork

The procedures for joining a mobile subnetwork are subnetwork specific and the failure will typically be the result of a configuration failure, failure in subnetwork equipment or erroneous protocol operation. A local log of the event for later analysis will usually be sufficient. However, an unusually high number of such events during a given reporting period may indicate a problem in a Ground Station that needs to be investigated and this should be reported to a Network Manager.

3.2 The ATN Router

The function of an ATN Router is, in principle, very simple: it is to forward CLNP packets received from an adjacent Router or End System to another adjacent Router or End System, such that that packet reaches its destination in the shortest possible time. With such a simple function, there are only six possible failure modes:

1. The router fails to forward the packet;

2. The packet is forwarded but with loss of data integrity;

3. The packet is forwarded but to a different router or End System than that which would have met the criteria for the packet reaching its destination in the shortest possible time.

4. The router suffers a complete system failure.

5. Improper use of a subnetwork.

6. The ATN also supporting roaming across mobile subnetworks and this introduces the failure mode of failure to complete the Mobile Route Initiation Procedures.

3.2.1 Failure Mode #7: Failure to Forward

This failure mode may be defined as the silent discard of a packet which a correctly functioning router should have forwarded.

The consequence of this failure is that the packet will naturally fail to get to its intended destination. This may be a transient or permanent failure mode and the packet may have been sent under the connection mode or connectionless transport service.

Under the connection mode transport protocol, the failure of a packet to reach its destination is detected by the non-receipt of an acknowledgement in the expected period. The transport layer is then required to re-send the packet. If the failure was transient, then the retransmission will succeed and the only detectable result of this is that the transit time for the packet will be longer than expected. If the failure is permanent then, even after the permitted number of retries, the packet will not be delivered and the transport connection will be declared as disconnected and reported as such to the application end user.

Under the connectionless transport protocol, the failure of a packet to reach its destination is not reported to the end user, regardless of whether it is a transient or permanent failure. If the application's delivery requirement is "at most once", then this is satisfactory as the requirement is being met. If the delivery requirement is "at least once" then application will itself need to detect this event and recover from failure to deliver by retransmission on no response (a typical feature of connectionless application protocols that have this delivery semantic). As with the connection mode transport protocol, recovery is made by retransmission. For transient failures this will be detectable as a longer than normal time to complete the transaction. For permanent failures this will result in a failure to perform the required transaction.

Transport and (for connectionless communications) application level notifications of an unacceptable degradation in the measured Quality of Service (e.g. by round trip delay computation) or complete loss of the transport connection are the possible trigger events.

3.2.2 Failure Mode #8: Loss of Data Integrity

In the ATN, both connection mode and connectionless transport protocols are required to support the checksum parameter. This is used to detect end-to-end loss of data integrity. Thus, if a Router fails to preserve the integrity of the packets passing through it, this will be detected by the receiving transport layer protocol, and the packet discarded.

The consequence of loss of data integrity is thus discard of the affected packet by the transport provider. In the connection mode case, this should be recoverable by retransmission. In the connectionless case, the impact on the application will be identical to the preceding failure mode.

A certain level of such corrupt packets should be tolerated. However, if the number of such discards exceeds some defined threshold then this is indicative of a more serious problem (e.g. a failing subnetwork or router) and needs to be reported to a Network Manager. The trigger event is thus transport provider (connection mode or connectionless) discard of a TPDU for checksum failures exceeding some reporting threshold during a given reporting period.

3.2.3 Failure Mode #9: Mis-routing

This failure mode by be defined as a incorrect or non-optimal routing decision.

As a result, a mis-routed packet may either reach its intended destination via a longer path than it should have done, or will be discarded because:

- the path it takes is longer than its permitted “lifetime” (e.g. the packet is in a routing loop), or
- because it is passed to a router that does not know how to route it.

Thus the immediate consequence of this failure mode is that either the affected packet will be delivered to its destination with a longer than expected transit delay, or that it will fail to be delivered.

In the latter case, the consequence is thus identical to failure mode #1. In the former case, this is only an issue if the transit delay is longer than required by operational requirements.

CLNP Packet discard is the main indication of this problem and should be identified as the trigger event. However, transport (and for connectionless communications application level) notifications of an unacceptable degradation in the measured Quality of Service (e.g. by round trip delay computation), or of a complete loss of the transport connection, are useful indicators as to the severity of the problem and can thus also be considered as trigger events. Indeed, as the packet discard report may be in a different Network Management Domain from the failing router and the affected end users, these may be critical trigger events from the point of view of initiating the fault diagnosis and repair procedures.

3.2.4 Failure Mode #10: Total System Failure

Total System Failure will result in a failure to forward any packets currently in transit through the router. The consequence of this is as discussed above in 3.2.1, and requires end-to-end recovery.

The impact of such a failure should be transient as it is the responsibility of neighbour routers to identify the failure of an adjacent router, to adjust their own routing tables as a result, and to report any consequential changes to other adjacent routers. These routers should notify the Network Manager that an adjacent router has failed, and this is the trigger event on which fault diagnosis and repair procedures will be initiated.

3.2.5 Failure Mode #11: Improper use of a subnetwork

This failure mode includes failure to establish required subnetwork connections in a timely manner, connections established using incorrect parameters, and incorrect use of subnetwork specific procedures and protocol.

In the latter two cases, the problem may be detected by the remote System. This will either be from an error report generated by the subnetwork (including a subnetwork connection call clearing reason), or by validation of call or packet parameters. However, not all such errors will be detectable (e.g. an incorrect throughput requirement).

Failures to establish a required subnetwork connection or path, or use of incorrect (but in range) parameters can result in there being a less than planned capacity in the network and, in extreme cases, loss of the end-to-end communications path. Alternatively, they could result in

over-capacity or in use of a more expensive path, and hence increased network cost. Lower than planned capacity will be visible to the end user as increased transit delay or loss of communications.

The Trigger events for this failure mode are thus:

- End System reports of unacceptable degradation in end-to-end transit delay or in unexpected transport connection loss, or
- ATN Router or End System reports of unexpected subnetwork connection loss or subnetwork connection parameter problems.

3.2.6 Failure Mode #12: Failure to Complete Mobile Route Initiation Procedures

Failure to complete the Mobile Route Initiation Procedures will typically be due to a configuration error or a protocol error. A local log of the event for later analysis will usually be sufficient. However, an unusually high number of such events during a given reporting period may indicate a problem in an Air/Ground Router that needs to be investigated and this should be reported to a Network Manager.

3.3 The End-to-end ATN Internet Service

The ATN Internet service, as provided to a user of the ATN Internet, is defined to be the OSI Transport Service, either connection mode or connectionless. The connection mode service provides for reliable stream mode communications, within the limits defined by a given Quality of Service, and with the semantic that each message sent over a transport connection will not be delivered until all messages sent before that message have been delivered. In an ATN Context, the transport service must also report a failure to maintain the expected Quality of Service, as application specific fallback procedures may need to be invoked in such cases.

The connectionless service is a simple datagram service with each message sent having a delivery probability of less than one and with no sequencing guarantees.

The connection mode failure modes are:

1. Failure to establish a Transport Connection
2. Mis-sequencing, Corruption, loss and/or mis-delivery of user messages;
3. Failure to report an unacceptable Quality of Service degradation including transport connection loss, and
4. Uncommanded loss of a Transport Connection.

The only failure mode of the connectionless transport service that can properly be called a failure mode is delivery of a message containing unreported errors (as in item 3 above). Failure to deliver cannot be regarded as an error as the service does not guarantee delivery.

3.3.1 Failure mode #13: Failure to establish a Transport Connection

The most likely reason for this failure mode is the lack of a network communications path, or the remote End System not being in an operational state. However, it can also result from

incorrect operation of either the initiating or responding transport provider, or the lower layer functions in either End System.

As the problem may be due to a network error affecting the end user, as discussed in 3.2.1, this event should be reported to a Network Manager and is a trigger for fault diagnosis and repair.

3.3.2 Failure Mode #14: Mis-sequencing, Corruption, loss and/or mis-delivery of User Messages

This failure mode is due to incorrect operation of the transport layer software. It can be due to incorrect assignment of TPDU sequence numbers, in packet re-ordering errors by the receiving transport entity, or in failure to detect and handle network errors.

In the ATN, Applications require a high availability, reliable transport service. As this failure mode demonstrates, this can only be provided if the transport layer software can be relied upon and hence, high quality design and implementation of the transport layer software and extensive testing are both required and assumed by this analysis.

As this is not an internet error but a software error in the End System, it is not useful to report such a problem to a Network Manager. However, such problems should be logged locally, if detected by an application, as well as reported to the End User if there are implications for the correct operation of the application.

3.3.3 Failure Mode #15: Failure to Report an Unacceptable Quality of Service Degradation

This failure mode will result from errors in the measurement of round trip delay by the sending transport protocol, or is due to a failure to detect/report non-receipt of an expected AK TPDU. Such a failure mode can only be detected by the application maintaining its own checks on round trip delay and “liveness” of the transport connection.

Again, this is not an internet error but a software error in the End System, and it is not useful to report such a problem to a Network Manager. However, such problems should be logged locally, if detected by an application, as well as reported to the End User if there are implications for the correct operation of the application.

3.3.4 Failure Mode #16: Uncommanded loss of a Transport Connection

A transport connection may be lost because of network errors, the failure of the remote End System, or because of internal protocol errors within the transport provider. This will be reported to the user who will need to take alternative action.

As the problem may be due to a network error affecting the end user, as discussed in 3.2.1, this event should be reported to a Network Manager and is a trigger for fault diagnosis and repair.

3.3.5 Summary

Transport failure modes can be due to underlying network problems or failures within the implementation of the transport provider. If the problem is unreported to the transport service user, detection requires duplication of transport protocol functionality within the application.

There is thus a fundamental difference between transport layer failures and failures in lower layers. Lower Layer failure modes, if visible to the end user, are detected and often recovered from by the independent operation of the end-to-end transport protocol. However, failures within

the transport protocol are neither detected nor recovered from unless detected by the user application detecting an inconsistency the in data transferred over the transport connection (e.g. an out of sequence message sequence number).

Whilst some applications may justify duplicate transport functionality, in general, the emphasis for fault management in the transport layer should be placed on fault prevention by high quality implementation and testing of the transport protocol. On the other hand, a much higher number of errors can be tolerated in the implementation of lower layer functions, because reliance can be placed on the transport protocol to detect and often recover from the fault.

4. Fault Management Procedures

This analysis has identified the following trigger events for Network Manager action:

1. Notification of Transport Connection Loss
2. Reports of Unacceptable degradation of the measured end-to-end Quality of Service (i.e. excessive transit delay).
3. CLNP packet discard reports,
4. Unexpected Subnetwork Connection Loss
5. CLNP Header Checksum Error Reports
6. SNDCF Deflate checksum error reports
7. Transport level checksum error reports.
8. Notification of adjacent Router failure.
9. Notification of failure to establish a transport connection.
10. Notification by an Application of errors in the provision of the Transport Service.
11. Failure to Join a Mobile Subnetwork
12. Failure to Complete Mobile Route Initiation Procedures
13. Subnetwork Connection Parameter Problem

The responses to the trigger events are considered in more detail below.

4.1 Response to Triggers

4.1.1 Trigger #1: Notification of Transport Connection Loss

A transport connection failure may imply a lower layer problem and thus needs to be reported to a Network Manager as well as the end user. However, it should be noted that there is probably no requirement to report transport connection failure on the airborne side as long as the ground side is responsible for reporting and responding to such an event.

A Network Manager will not necessarily respond to every report of a lost transport connection. In many cases these will be due to transient problems (e.g. an aircraft going out of range of a VHF transmitter). However, when the rate of such events exceeds some specified limit, the cause of such events will need to be investigated in order to determine any common mode problem.

To investigate this, the Network Manager will need to trace out the route (or at least common parts of the affected routes) in order to try and identify the point of failure. The CLNP Echo PDU may be used for this purpose.

The Network Manager will identify a suitable starting point in the network and from this point, initiate an Echo Request, with route tracing enabled, and with the Error Report flag set. The destination of the echo will be the identified end of the problematic communications path. A series of such Echo PDUs will be sent over some period. The Security Options and priority should reflect that of the data for which problems have been found.

The Network Manager will then await the responses:

- All being well, an Echo Response will be received containing the Echo request as received at the point of echo; this will include the route trace information.

- If a problem is encountered along the way, then an Error PDU should be returned containing the Echo request and route trace information at the point at which the problem was encountered.

- In the worse case, no response will be received, and the Network Manager will need to probe the route with “Echoes” directed to different systems along the expected route until the “black hole” is found.

If no problems are reported for any of the Echoes, then the problem was transient and has gone away, or the problem exists in a different part of the network and the above needs to be repeated for alternative network paths.

If a problem is reported then the problem router or data link will be the next “hop” in the path from which the problem was reported, and this will need to be investigated further. It is quiet possible that this will be in a different Network Management Domain to the originator of the Echo Request and that liaison between the responsible Network Managers will be required to repair the fault.

Analysis of faulty routers and subnetworks is dealt with below in 4.1.12.

4.1.2 Trigger #2: Notification of end-to-end Quality of Service Degradation

A transport layer using a connectionless network service can only realistically measure TPDU checksum failures across all transport connections, and round trip delay (which can be assumed to be the sum of the transit delay in each direction). As TPDU checksum failures result in retransmissions and hence increased transit delay, a notification of end-to-end Quality of Service Degradation equates to a notification of an unacceptable end-to-end transit delay.

As with a complete loss of a transport connection, a notification of an unacceptable end-to-end transit delay may imply a serious problem in the Internet and thus needs to be reported and investigated. The investigation will be as above - connection loss is really just an extreme case of QoS degradation. Similarly, there is no requirement to report the problem over an air/ground data link.

4.1.3 Trigger #3: CLNP Packet Discard due to Lifetime Expiry or a Routing Problem

Packet discard events for packet lifetime expiry or a routing problem, are indicative of a network problem either transient or permanent. They are reported by the Router that identifies the problem, but which is unlikely to be the source of the problem. As with transport connection loss events, a single discard event is unlikely to warrant investigation. However, a high rate of such reports, especially if concentrated in a single part of the Internet, will require investigation.

It should also be noted that some routers will experience a high number of certain types of discard. For example, a “Backbone Router” will regularly discard unroutable packets to aircraft that are currently “out of contact”. It follows that notification thresholds will differ by discard reason. Alternatively, such an issue could be avoided if such unroutable packets were always directed to a specially designed “sink” router, whose only role was to discard such packets.

Packet discard events may also be reported over the air/ground data link to the Network Manager for the Air/Ground Service provider. This is necessary in order to report mis-routing over the air/ground data link and to avoid extended periods of denial of service to affected aircraft. However, such a requirement could be avoided if uplinked packets were always sent with an “error reported requested” in event of discard, and such error reports, on their return, were reported to an appropriate Network Manager.

The first step in the analysis of the problem, will be inspection of the pattern of packet discards and the reason for discard. In itself, this may be sufficient to identify a likely source of the problem. Alternatively, an area of the Internet (comprising a number of Routers and Data Links) may be suspected, and probed. As above, CLNP Echo PDUs are the most effective mechanism for probing the suspected Routers and Data Links and a number of such probes will need to be may over various communications paths in the area. This should identify the source of the problem, with further analysis of the failed router and corrective action discussed in 4.2.1 below.

4.1.4 Trigger #4: Unexpected Subnetwork Connection Loss

Data link loss should be reported by the affected Routers. In most cases, the repair procedure will be network specific and will require action by the service provider, as discussed in 4.2.1. However, the reason may be due to an incorrect use of the subnetwork by either the reporting system, or its remote peer. To determine if this is the case will require analysis of logs at each end and, in particular, the parameters used to establish the subnetwork connection.

As the report is provided by both ends of a failed data link, there is no need to report this problem over an air/ground data link, as the ground side should always be able to log or report the problem to a Network Manager.

4.1.5 Trigger #5: CLNP Checksum Failure

CLNP Checksum Failure is detected and reported by a Router when it discards the affected packet, and is indicative of a problem on the data link or the router from which the packet was received. A single such event can probably be ignored. However, once the rate goes above a given threshold or if several routers in the same area start reporting similar problems, then investigation will be necessary.

This event will need to be reported over air/ground data links as problems in the Ground Station part of the data link may otherwise go unreported and un-repaired. Unless, as discussed in 4.1.3 above, error reports are enabled for uplinked packets.

Both the connecting data link(s) and the other router will need to be investigated. Inspection of data link statistics may be diagnostic, as discussed in 4.2.1. However, probing the router and data link with CLNP Echo PDUs, as discussed below in 4.1.7 may also be necessary, in order to determine which is at fault.

4.1.6 Trigger #6: SNDCF Deflate checksum error reports

SNDCF Deflate checksum error reports are indicative of a problem on the data link or the router from which the packet was received. As above, a single such event can be ignored. However,

the problem needs to be investigated once the rate of such errors exceeds a given threshold. The fault diagnosis and repair procedure are as for 4.1.5 above.

As the report may be provided by both ends of the data link, there should be no need to report this problem over an air/ground data link, as the ground side should report the problem to a Network Manager.

4.1.7 Trigger #7: Transport level checksum error reports

The source of this problem can be an otherwise undetected network fault or a problem in the sending or receiving End System. As with other such reports, the report is only of interest when the number of such events exceeds a given threshold during a reporting period.

As regards Air/Ground Communications, the problem may be asymmetric (e.g. introduced on an uplink only by a faulty Ground Station) and therefore may have to be reported over an air/ground data link as otherwise such faults will continue to be unreported and un-repaired. However, if the Deflate compression is used, this requirement should be avoidable, as the Deflate compression scheme and checksums should detect the problem such that it is recognised by the Air/Ground Router.

In general, the diagnosis of such a problem will require a more sophisticated use of ECHO PDUs than is needed for detecting mis-routing problems. A series of ECHO PDUs will need to be despatched over the suspected communications path with known test data patterns as their payload. The ECHO responses should contain the same test data reflected back, and comparison of the reflected test data with what was originally sent will reveal any errors that have been induced on the route.

Once errors have been identified on a communications path, performing the same function on smaller segments of the path will reveal the router or data link that is inducing the error. Further analysis of the failed component is discussed below in 4.1.12.

4.1.8 Trigger #8: Notification of adjacent Router failure.

This error report will be due to either an actual failure in the adjacent router or the failure of the data link supporting the adjacency. Failure of an actual router should be clear because the failure should be reported by all adjacent routers. A failure report from only one router will typically indicate a failed data link, and should usually be accompanied by a subnetwork connection failure report (for a connection mode subnetwork) and a report from the router regarded as having failed reporting the apparent failure of the first router!

There should be no requirement to report this problem over an air/ground data link. Firstly, such reports will be commonplace over mobile subnetworks and secondly, a failure in an Air/Ground Router will be reported by other Ground Routers.

If a Router has really failed then repair will usually require physical intervention either to restart the router in the event of a software crash, or a hardware repair.

4.1.9 Trigger #9: Notification of failure to establish a transport connection

The underlying reasons should be no different to those that result in connection loss and diagnosis and repair should be as in 4.1.1.

4.1.10 Trigger #10: Application Reported Transport Service Problem

If an application detects incorrect operation of the transport service, this should be reported to a Network Manager as the fault may lie in the underlying transport protocol. Investigation of such a problem is discussed in 4.2.3.

4.1.11 Trigger #11: Failure to Join a Mobile Subnetwork

If the number of failures exceeds a specified threshold then a Ground Station failure should be suspected. Logs of individual events will need to be inspected in order to find a common mode problem. However, the analysis of the problem will be specific to each subnetwork type.

4.1.12 Trigger #12: Failure to Complete Mobile Route Initiation Procedures

If the number of failures exceeds a specified threshold then an Air/Ground Router failure should be suspected. This could be in the operation of the Mobile SNDCF, ES-IS or IDRP protocols, all of which are involved in this procedure. Logs of individual events will need to inspected in order to find a common mode problem (e.g. in Router configuration).

If no common problem can be found through such an inspection, then a software fault should be suspected. The Router should be isolated and restarted, with a “snapshot” of the system’s state taken for post-mortem analysis.

4.1.13 Trigger #13: Subnetwork Connection Parameter Problem

The event is generated by an ATN Router or End System that detects the incorrect use of a subnetwork service by its remote peer. The most likely reason for this is a configuration error in the remote system. This needs to be investigated and corrected.

4.2 Component Fault Analysis

4.2.1 Analysis of a possibly faulty Data Link

Analysis of a suspected data link problem will always have a dependency on the network type and the test tools will be specific to that subnetwork. For example, testing of a faulty coaxial cable Ethernet will require a Reflectometer in order to find the source of the fault.

In general, a suspected data link will require analysis of the data link statistics gathered by the routers at each end and probing with test data and data link specific diagnostics. Inspection of error related statistics will often reveal problems, and comparison of sent and received data counts at each end of a data link will identify mis-routing and data loss problems in the subnetwork.

It should be noted that such an analysis does not have to wait for a failure to be reported and can be conducted regularly as part of a preventative maintenance schedule, often saving considerable time in diagnosis procedures and avoiding altogether end user visible problems.

A failed subnetwork will normally need to be taken out of service, until it is repaired by the service provider, and the mechanisms need to be provided in order to isolate a subnetwork during repair and to bring it back into service later on. This will require the availability of suitable Systems Management “Actions” in attached End Systems and Routers.

4.2.2 Analysis of a possibly faulty Router

A Router can be faulty either in its routing function or in its use of subnetworks.

4.2.2.1 Analysis of a Router's Routing Function

A suspected Router will need to be tested by both external probes i.e. Echo PDUs deliberately sent through it on all possible routes (including all variations of ATSC Class, etc.) that pass through the router – this is in order to test the proper operation of the forwarding function – and through inspection of internal data structures.

The data structures that need to be checked are:

1. The Forwarding Information Base (FIB), and

2. The Routing Information Base (RIB).

The FIB is used to determine the forwarding decision for each packet and, by definition, erroneous forwarding decisions will always manifest themselves in the FIB. Identification of an incorrect FIB entry(ies) will be diagnostic of a failed router. Such an entry should also be traced back to the RIB.

The RIB contains the routing information on which the FIB is constructed. It is learnt from information received from other routers and from local information sources. It is also the source of routing information passed to other routers. In an operational ATN Router, elements of the RIB may be the responsibility of three different Routing Information Exchange Protocols: IDRP, IS-IS and ES-IS.

A lack of consistency between the erroneous FIB entries and the RIB will indicate a software error in the synthesis of the FIB from the RIB. Otherwise, the problem will be in the Routing Information exchanged between Routers. This could mean that the faulty router is really an adjacent router passing incorrect routing information to this router. Determination of the actual faulty router will require inspection of the RIBs of adjacent routers as well as the RIB of the suspected router, looking for inconsistencies between the RIBs in the different routers and with the actual network topology. It is possible that the search for the faulty router may even lead to routers more than one hop away.

Once an inconsistency has been found, then the actual source of the problem can be identified. An inconsistency between RIBs in adjacent routers will point to errors in the exchange of routing information. In the case of IDRP, this could be policy inconsistencies. An inconsistency between a RIB and the known network topology will point to a software defect in that router.

Repair of the faulty router may be achieved by:

1.	System isolation and restart.	This is probably the safest strategy as it returns the system to a known state, but does mean that the router is unavailable for a short period. It is probably the only sensible strategy for a software error. Note that in the case of a software problem, a "snapshot" of the system state should be taken for post-mortem analysis, so that the bug can be identified and fixed.
2.	RIB Refresh	This is an IDRP procedure that forces the exchange of all current routes between a pair of adjacent routers. It is a valid repair strategy when an inconsistency is found

		between the adj-RIN-in and adj-RIB-out of a pair of adjacent BISs – possibly due to undetected errors in underlying communications path.
3.	Forced transmission of Link State PDUs	This is applicable to IS-IS and can be used to force early re-synchronisation between the RIBs of adjacent routers when an inconsistency is detected.

It should be noted that FIB/RIB and RIB/RIB consistency checks do not need to wait for a suspected error. Under a preventative maintenance schedule, such checks can be performed regularly and used to detect and repair faults before they become problems. This can also save considerable time in diagnostic effort and help to maintain a high operation availability. Indeed, one of the most serious potential problems facing a fault tolerant network is that a previously unknown fault is found in the backup route when a primary route fails, this leading to a total loss of the affected communications path. Regular FIB/RIB and RIB/RIB consistency checks can avoid this problem occurring.

With consistency checks, a “false positive” may result from real time problems - for example, if a consistency check is performed during a RIB update. However, it is unlikely for there to be a need to suspend such updates during consistency checks, as repeating the consistency check will eliminate these “false positives” which are always short lived.

4.2.2.2 Analysis of a Router’s Use of Subnetworks

The result of incorrect use of Subnetworks will be subnetwork connections which should be there, but aren’t, subnetwork connections with too much or too little capacity, or subnetwork connections which have been established but which are not being used.

In order to investigate these possible problems, a Network Manager must be able to inspect the configuration and state information associated with each subnetwork in order to determine the subnetwork(s) and subnetwork connections in use, the parameters with which each subnetwork connection was established, and the send receive and error counts on each subnetwork.

The Network Manager will then need to determine if the pattern of available subnetwork paths and subnetwork connection capacities (and other operational parameters) is in line with expectations, and that the usage of each subnetwork path is also in line with expectations. If this is not the case, then the Network Manager will need to determine why.

The most likely reason is a configuration or policy error. Otherwise, an internal software fault will need to be suspected. In the former case, repair is by correction of the error. In the latter case, System isolation and restart should correct the problem in the short term. A software bug fix will be required for a long term solution.

It should be noted that configuration faults that result in excess capacity will not result in externally visible problems except, perhaps, when the bill arrives! To detect such problems will require regular inspection of a system’s, and especially a router’s use of subnetworks. A regular inspection is also useful to detect similar problems that can result in externally visible faults.

4.2.3 Analysis of a possibly faulty End System

Faults may be suspected in the subnetwork access, network and transport layer components of an End System. These problems may be difficult to diagnose remotely as the fault may affect the system’s ability to communicate with other systems, and local diagnostic tools will thus always have to be provided by an End Systems supplier.

However, remote diagnosis may be appropriate for many suspected faults. Those in the subnetwork access components can be identified by comparison of sent and received data counts, etc. as discussed above in 4.2.1. Faults in the network components can be identified using similar procedures to those for a Router. An End System will have both FIB and RIB type structures, although these may be very simple, with the FIB having often only one entry, and never more than a “handful” of entries, and the RIB holding ES-IS related information only. However, faults may still be due to inconsistencies between these structures and the actual network topology and should still be investigated.

Transport level faults will require diagnose techniques additional to those required for Routers and Subnetworks, A Test Application will be needed to diagnose suspected transport service problems. Such a Test Application will need to be able to test and verify all transport service functions and to act as test data generator and data “reflector” in order to test for data integrity problems. The typical scenario for such a Test Application, is with the End System isolated from the network. However, there may be circumstances when it needs to be tested online.

Repair of a faulty End System will typically require isolation and restart for software problems and the intervention of an Engineer for hardware problems. As with router software problems, a system “snapshot” should be taken before a restart in order to provide for a later post-mortem analysis in order to find the bug.

4.3 Derived Trigger Responses

The preceding analysis has identified three cases where analysis is needed to detect a fault and where preventative maintenance is appropriate by regular inspection on Systems Management information. These can be performed by automatic processes raising an event report when a problem is found. These are also triggers for Network Manager action and described here as “Derived Triggers”. This is because the event is not raised directly by a network component, but by the Network Management Station as a result of its monitoring function.

4.3.1 Trigger #14: Suspected Data Link Problem

This trigger arises out of a regular monitoring of error counts of a given data link, and comparison of send and receive data counts at each end of a data link. An excessive number of errors, or a significant enough discrepancy in the counts, during some reporting period, will cause the event to be raised. Further network specific diagnostic procedures may then be appropriate.

There is probably a need to perform this procedure over an air/ground data link, at least on an occasional basis. Again, this is so problems in Ground Stations can be detected early. However, the use of the Deflate compression algorithm should avoid this requirement. When Deflate is used, any errors in uplinked or downlinked packets, or missing packets will be recognised by both Airborne and Air/Ground Router. The Air/Ground Router should monitor such incidents and report an increase from the anticipated undetected error rate for the air/ground data link.

The repair procedure will again be network specific and will require action by the service provider. The data link may need to be taken out of service until the problem is fixed.

4.3.2 Trigger #15: RIB/FIB and RIB/RIB Inconsistencies

Regular comparison of RIB and FIB information in the same Router, and of RIB information in adjacent routers may reveal inconsistencies. Such inconsistencies are a trigger event for Network Manager action. This is also another case, where co-operation is needed across Management Domains.

There is probably a need to perform this procedure over an air/ground data link, at least on an occasional basis. Again, this is so problems in Ground Stations can be detected early. However, to avoid air/ground communications overhead, an alternative strategy may need to be adopted, with a logging of routing events and FIB updates on both sides of an Air/Ground datalink for offline comparison later and on a sample basis.

The handling of such inconsistencies has been dealt with in 4.2.2.1.

4.3.3 Trigger #16: Improper Use of Subnetworks

Regular checks on the set of available subnetwork paths, the operational parameters and usage of each path can reveal configuration and policy errors, or software defects, that can lead to end-to-end communications problems or unnecessary costs.

The handling of such inconsistencies has been dealt with in 4.2.2.2.

5. Requirements on the Management Information Base

5.1 System Level

1. **[REQ 1]** ATN Routers are required to provide for remote restart on command from a Network Manager.

2. **[REQ 2]** Remote isolation of ATN Routers is required. That is graceful termination of the operational state and entry into a state where the router only responds to Systems Management requests.

5.2 Applications

[REQ 3] Applications that implement compliancy checks on the operation of the Transport Service are required to notify a Network Manager of incorrect operation of the Transport Service.

5.3 Transport Layer

1. **[REQ 4]** The Transport Provider is required to notify a Network Manager when a Transport Connection fails to be established.

2. **[REQ 5]** The Transport Provider is required to notify a Network Manager when a Transport Connection is lost.

3. **[REQ 6]** The Transport Layer is required to notify a Network Manager when the end-to-end transit delay as measured by the transport provider and derived from the round trip delay, falls below a specified threshold (typically application specific - e.g. based on ATSC Class).

4. **[REQ 7]** The Transport Layer is required to notify a Network Manager when the number of TPDU discards due to checksum validation failure exceeds a Network Manager specified threshold during a given reporting period.

5.4 CLNP

1. **[REQ 8]** The connectionless network service provider is required to notify a Network Manager when the number of CLNP PDUs discards for Header Checksum verification, lifetime expiry or routing problems exceeds a specified threshold. Such thresholds will need to be specified by discard reason.

2. **[REQ 9]** The connectionless network service provider is required to notify a Network Manager when an ECHO Request or Response is received and to provide information contained in the PDU.

3. **[REQ 10]** The connectionless network service provider is required to notify a Network Manager when an Error PDU is received at its addressed destination, and to provide information contained in the PDU.

4. **[REQ 11]** The connectionless network service provider is required to support the remote invocation of the Echo Request function and option selection.

5. **[REQ 12]** The connectionless network service provider is required to provide remote access to its Forwarding Information Base (FIB).

5.5 IS-SME

1. **[REQ 13]** The IS-SME is required to report when the number of times a failure to complete the Route Initiation procedures event is logged exceeds a Network Manager specified threshold during a given reporting period.

5.6 IDRP

1. **[REQ 14]** IDRP is required to provide remote access to its Routing Information Base (RIB), permitting the download of the complete RIB by a Network Manager.

2. **[REQ 15]** IDRP is required to report loss of an adjacency.

3. **[REQ 16]** IDRP is required to support remote start up and shutdown of Ground/Ground adjacencies with other Routers (for fault isolation).

4. **[REQ 17]** IDRP is required to support remote invocation of the RIB Refresh procedure.

5.7 IS-IS

1. **[REQ 18]** IS-IS is required to provide remote access to its Routing Information Base (RIB), permitting the download of the complete RIB by a Network Manager.

2. **[REQ 19]** IS-IS is required to report loss of an adjacency.

5.8 ES-IS

1. **[REQ 20]** ES-IS is required to provide remote access to its Routing Information Base (RIB), permitting the download of the complete RIB by a Network Manager.

5.9 SNDCF

1. **[REQ 21]** The Mobile SNDCF is required to report packet level Deflate checksum failures (Air/Ground Routers only).

2. **[REQ 22]** Connection Mode SNDCFs are required to report unexpected subnetwork connection loss and subnetwork reset.

3. **[REQ 23]** SNDCFs are required to support remote start up and shutdown of data links.

5.10 Subnetworks

1. **[REQ 24]** Subnetworks are required to keep counts of packets sent and received, and of error counts where applicable, and to provide remote access to such statistics.

2. **[REQ 25]** Subnetworks are required to report error counts that exceed a specified threshold during a set reporting period.

3. **[REQ 26]** Subnetworks are required to provide subnetwork specific diagnostics and test procedures, as appropriate for the subnetwork type, and to support their remote use by a Network Manager.

4. **[REQ 27]** Subnetworks are required to report the set of existing subnetwork connections and the operational parameters for each such connection.

5. **[REQ 28]** Mobile Subnetworks are required to report when the number of times a failure to join a mobile subnetwork event is logged exceeds a Network Manager specified threshold during a given reporting period.

6. Requirements on Implementations

6.1 End Systems

1. **[REQ 29]** Transport layer implementations are required to support the measurement of round trip delay on a per transport connection basis, and hence to estimate the end-to-end transit delay. An unacceptably long transit delay is to be reported both to the service user and (for Ground Systems) a Network Manager.

2. **[REQ 30]** CLNP Header checksums are required, in order to detect subnetwork problems at source.

3. **[REQ 31]** A CLNP Error Report shall be requested for all Data PDUs addressed to airborne destinations

6.2 ATN Routers

1. **[REQ 32]** The Deflate compression algorithm is required on all air/ground data links in order to support early detection of Ground Station problems.

7. Requirements for Network Manager Tools and Procedures

In addition to reporting of Notifications and access to Managed Objects, a Network Management Station (NMS) needs to support the following functions for fault management:

1. **[REQ 33]** A Network Manager shall be provided with a means to illustrate the topological distribution of packet discard and transport connection loss/QoS degradation events

2. **[REQ 34]** A Network Manager shall be provided with tools to generate ECHO PDUs from any point in the Internet, and to correlate the notifications of ECHO responses and Error Responses with ECHO requests. The generation and use of Test Data patterns as the payload of ECHO Request PDUs shall also be supported, including monitoring for inconsistencies between responses and requests.

3. **[REQ 35]** The NMS shall regularly monitor sent and received data counts on subnetwork connections, and report variations between those at each end of a data link that exceed a set threshold during a reporting period.

4. **[REQ 36]** The NMS shall regularly monitor FIB and RIB state on each Router and report inconsistencies between a Router's FIB and RIB state.

5. **[REQ 37]** The NMS shall also check the consistency of RIBs in adjacent Routers and report inconsistencies.

6. **[REQ 38]** The NMS shall check the availability of subnetwork paths for each End System and Router, the Operational Parameters and usage of each such path, against that expected, and report significant deviations.

7. **[REQ 39]** There is a requirement for the exchange of Management Information between Management Domains in respect of:

    - Subnetwork send, receive and error counts
    - RIB/FIB Information in Boundary Routers
    - CLNP ECHO Request Received Notifications
    - CLNP Error Report Received Notifications

Appendix A: Failure Mode to Triggers Cross-Reference Table

Table A-7-1Failure Mode to Trigger Cross-Reference

	Trigger # 1	Trigger # 2	Trigger # 3	Trigger # 4	Trigger # 5	Trigger # 6	Trigger # 7	Trigger # 8	Trigger # 9	Trigger # 10	Trigger # 11	Trigger # 12	Trigger # 13	Trigger # 14	Trigger # 15	Trigger # 16
Failure Mode #1				✓												
Failure Mode #2														✓		
Failure Mode #3			✓			✓	✓									
Failure Mode #4														✓		
Failure Mode #5			✓			✓	✓									
Failure Mode #6											✓					
Failure Mode #7	✓									✓					✓	
Failure Mode #8							✓									
Failure Mode #9	✓	✓	✓												✓	
Failure Mode #10								✓								
Failure Mode #11				✓									✓			✓
Failure Mode #12												✓				
Failure Mode #13									✓							
Failure Mode #14																
Failure Mode #15																
Failure Mode #16	✓															

Table A-7-1 provides a cross-reference between the identified failure modes and the identified trigger events. It should be noted that some failure modes do not result in trigger events. These are failures internal to the transport service on an End System and are logged locally rather than reported to a Network Manager. Trigger #14 is a derived trigger event and does not correspond to any explicit failure mode.

Appendix B: Triggers to Requirements Cross-Reference Table

Table B-1 Triggers to Requirements

	Trigger																
	#1	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**	**#11**	**#12**	**#13**	**#14**	**#15**	**#16**	**Requirement**
REQ 1	✓	✓	✓				✓	✓	✓	✓		✓	✓		✓	✓	ATN Routers are required to provide for remote restart on command from a Network Manager
REQ 2	✓	✓	✓				✓	✓	✓	✓		✓	✓		✓	✓	Remote isolation of ATN Routers is required. That is graceful termination of the operational state and entry into a state where the router only responds to Systems Management requests.
REQ 3										✓							Applications that implement compliancy checks on the operation of the Transport Service are required to notify a Network Manager of incorrect operation of the Transport Service.
REQ 4									✓								The Transport Provider is required to notify a Network Manager when a Transport Connection fails to be established
REQ 5	✓																The Transport Provider is required to notify a Network Manager when a Transport Connection is lost.
REQ 6		✓															The Transport Layer is required to notify a Network Manager when the end-to-end transit delay as measured by the transport provider and derived from the round trip delay, falls below a specified threshold (typically application specific - e.g. based on ATSC Class).
REQ 7							✓										The Transport Layer is required to notify a Network Manager when the number of TPDU discards due to checksum validation failure exceeds a Network Manager specified threshold during a given reporting period.

	Trigger																
	#1	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**	**#11**	**#12**	**#13**	**#14**	**#15**	**#16**	**Requirement**
REQ 8			✓		✓												The connectionless network service provider is required to notify a Network Manager when the number of CLNP PDUs discards for Header Checksum verification, lifetime expiry or routing problems exceeds a specified threshold. Such thresholds will need to be specified by discard reason.
REQ 9	✓	✓	✓		✓		✓		✓					✓	✓		The connectionless network service provider is required to notify a Network Manager when an ECHO Request or Response is received and to provide information contained in the PDU.
REQ 10	✓	✓	✓		✓		✓		✓					✓	✓		The connectionless network service provider is required to notify a Network Manager when an Error PDU is received at its addressed destination, and to provide information contained in the PDU
REQ 11	✓	✓	✓		✓		✓		✓					✓	✓		The connectionless network service provider is required to support the remote invocation of the Echo Request function and option selection
REQ 12	✓	✓	✓						✓						✓		The connectionless network service provider is required to provide remote access to its Forwarding Information Base (FIB).
REQ 13												✓					The IS-SME is required to report when the number of times a failure to complete the Route Initiation procedures event is logged exceeds a Network Manager specified threshold during a given reporting period.
REQ 14								✓							✓		IDRP is required to provide remote access to its Routing Information Base (RIB), permitting the download of the complete RIB by a Network Manager

	Trigger																
	#1	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**	**#11**	**#12**	**#13**	**#14**	**#15**	**#16**	**Requirement**
REQ 15																	IDRP is required to report loss of an adjacency
REQ 16															✓		IDRP is required to support remote start up and shutdown of Ground/Ground adjacencies with other Routers (for fault isolation).
REQ 17															✓		IDRP is required to support remote invocation of the RIB Refresh procedure
REQ 18															✓		IS-IS is required to provide remote access to its Routing Information Base (RIB), permitting the download of the complete RIB by a Network Manager.
REQ 19								✓									IS-IS is required to report loss of an adjacency
REQ 20															✓		ES-IS is required to provide remote access to its Routing Information Base (RIB), permitting the download of the complete RIB by a Network Manager.
REQ 21						✓	✓							✓			The Mobile SNDCF is required to report packet level Deflate checksum failures (Air/Ground Routers only).
REQ 22				✓									✓				Connection Mode SNDCFs are required to report unexpected subnetwork connection loss and subnetwork reset
REQ 23														✓			SNDCFs are required to support remote start up and shutdown of data links
REQ 24														✓		✓	Subnetworks are required to keep counts of packets sent and received, and of error counts where applicable, and to provide remote access to such statistics

	Trigger																
	#1	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**	**#11**	**#12**	**#13**	**#14**	**#15**	**#16**	**Requirement**
REQ 25														✓		✓	Subnetworks are required to report error counts that exceed a specified threshold during a set reporting period
REQ 26														✓			Subnetworks are required to provide subnetwork specific diagnostics and test procedures, as appropriate for the subnetwork type, and to support their remote use by a Network Manager
REQ 27				✓									✓			✓	Subnetworks are required to report the set of existing subnetwork connections and the operational parameters for each such connection.
REQ 28											✓						Mobile Subnetworks are required to report when the number of times a failure to join a mobile subnetwork event is logged exceeds a Network Manager specified threshold during a given reporting period.
REQ 29		✓															Transport layer implementations are required to support the measurement of round trip delay on a per transport connection basis. An unacceptably long transit delay is to be reported both to the service user and (for Ground Systems) a Network Manager
REQ 30					✓												CLNP Header checksums are required, in order to detect subnetwork problems at source.
REQ 31			✓		✓												A CLNP Error Report shall be requested for all Data PDUs addressed to airborne destinations
REQ 32							✓							✓			The Deflate compression algorithm is required on all air/ground data links in order to support early detection of Ground Station problems

	Trigger																
	#1	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**	**#11**	**#12**	**#13**	**#14**	**#15**	**#16**	**Requirement**
REQ 33	✓	✓	✓		✓												A Network Manager shall be provided with a means to illustrate the topological distribution of packet discard and transport connection loss/QoS degradation events
REQ 34	✓	✓	✓		✓		✓		✓					✓	✓		A Network Manager shall be provided with tools to generate ECHO PDUs from any point in the Internet, and to correlate the notifications of ECHO responses and Error Responses with ECHO requests. The generation and use of Test Data patterns as the payload of ECHO Request PDUs shall also be supported, including monitoring for inconsistencies between responses and requests
REQ 35														✓			The NMS shall regularly monitor sent and received data counts on subnetwork connections, and report variations between those at each end of a data link that exceed a set threshold during a reporting period
REQ 36															✓		The NMS shall regularly monitor FIB and RIB state on each Router and report inconsistencies between a Router's FIB and RIB state.
REQ 37															✓		The NMS shall also check the consistency of RIBs in adjacent Routers and report inconsistencies
REQ 38																✓	The NMS shall check the availability of subnetwork paths for each End System and Router, the Operational Parameters and usage of each such path, against that expected, and report significant deviations

	Trigger																
	#1	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**	**#11**	**#12**	**#13**	**#14**	**#15**	**#16**	**Requirement**
REQ 39	✓	✓	✓		✓		✓		✓					✓	✓		There is a requirement for the exchange of Management Information between Management Domains in respect of: · Subnetwork send, receive and error counts · RIB/FIB Information in Boundary Routers · CLNP ECHO Request Received Notifications · CLNP Error Report Received Notifications