EUROCONTROL

ATNP/WG2/
WP168/
06 October 1995

AERONAUTICAL TELECOMMUNICATIONS NETWORK PANEL

WORKING GROUP TWO

Banff 9.10.95-13.10.95

Defects found in IDRP and Consequential Changes Required to the draft ATN Internet SARPs

Presented By Henk Hof

Prepared by Tony Whyman

SUMMARY

Validation of the draft ATN SARPs is continuing through both simulations and implementation testing. This work has discovered a number of defects in ISO 10747 which need to be countered by changes to the ATN draft SARPs. This paper presents a summary of the defects found so far, and recommends those changes necessary to the draft SARPs.

DOCUMENT CONTROL LOG

SECTION	DATE	REV. NO.	REASON FOR CHANGE OR REFERENCE TO CHANGE
	28-Sep-95	Issue 1.0	

TABLE OF CONTENTS

1. Introduction

Validation work on IDRP is currently underway, with interoperability testing being conducted between several implementations and initial simulation results also becoming available. This work has now uncovered several defects which need to be corrected in the ISO standard and taken into account when developing ATN equipment. This paper lists defects found so far.

2. ISO 10747 - Defects Found

2.1 Specification of Type 1 Authentication Algorithm

2.1.1 Problem Summary

This defect was discovered when an interoperability problem was found between two implementations. i.e. they refused to talk to each other because their authentication checksums never match.

The ISO Standard (annex B) specifies the type 1 authentication algorithm as being “based on” the MD4 message digest algorithm, as specified by RFC 1186, but did not elaborate on any differences there might be. Analysis of the interoperability problem showed that one implementation simply implemented MD4 as its message digest, straight out of the RFC, while the other implemented it exactly as ISO 10747 specified the message digest.

The former implementation took a little endian approach to (e.g.) the length field, as is mandated by the RFC "where each consecutive group of four bytes is interpreted as a word with the low order (least significant) byte given first". The other implementation took 10747 literally. There is nothing in annex B on byte order, and so the implementors referred to section 6 which specifies a big endian approach "when consecutive octets are used to represent a number, the lower octet number has the most significant value".

The second implementation does arguably implement the standard, although it is not believed that there was any intention to deviate from standard implementations of MD4. Discussion with the ISO editor has revealed that this was not a deliberate decision. The term “based on” was chosen in case MD4 was later changed and to avoid invalidating existing implementations of ISO 10747 without an explicit decision in ISO to follow any such change.

Firstly, it is clear that because of existing familiarity with MD4, the ISO specification is open to mis-understanding. On the other hand, when read literally it is self-consistent, although the unintended deviation from MD4 has resulted in a loss of the existing test suites for MD4 and may weaken the effectiveness of the algorithm as some of the analysis of its operation may be invalidated.

2.1.2 Proposed Resolution

It is believed appropriate to resolve this problem by change to the ISO standard, so that annex B implements MD4 as it is specified today. This will bring the standard into line with most implementations and ensure that the existing test suites and experience with MD4 are available to ISO 10747 implementations.

2.2 Reflecting Routes Back to the Originator

2.2.1 Problem Summary

There appears to be a missing requirement in ISO 10747, and one that is so obvious that you assume that it is there. The requirement is "do not advertise back to a BIS, a route that you have received from it". This can be generalised to "do not advertise to a BIS a route which contains its RDI in its path, or the RDI of a confederation which will be entered by a route advertised to that BIS".

If you do advertise such a route, then it can be argued that ISO 10747 clause 7.16 requires that it is silently discarded, and so is not an error and merely expensive. It is only an error when 7.16 fails to be applied and a receiving BIS detects this according to clause 7.20.3. However, as to why there is no such requirement is not known.

Perhaps, there was an argument that said that there should be no statement on this as (a) the decision on which routes to advertise is always a matter for routing policy and (b) to permit simple low cost implementations that did not scan the full RD_Path there should be no such requirement.

2.2.2 Problem Resolution

At the very least a note is needed warning of the need for a Routing Policy that avoids reflecting back routes. However, there are circumstances where permitting routes to be advertised to an RD through which they have already passed through can cause problems with route withdrawal, and this is discussed below.

2.3 Persistent False Routes

2.3.1 Problem Summary

This problem only manifests itself in configurations that permit routes to be available via alternate paths. Such a configuration is illustrated in Figure 1. This illustrates a configuration of four BISs, which are fully interconnected.

At some time a new route “R” is received by router A and distributed to the other routers, which further distribute the route amongst themselves. In the figure, a notation is adopted to identify each distributed route from its RD_Path information. The notation is R(x,y,...), where the n-tuple (*x,y,..).* represents the path that the route has followed, by identifying each Router that the route has passed through, and in the order in which the routers have handled the route.

As Figure 1 illustrates, the initial distribution of route “R” is assumed to be straightforward, with each router choosing its preferred route on the basis of lowest hop count. Routers B,C,D therefore select the route they have received directly from A as their preferred route and re-advertise it to the other routers. Although not illustrated for the sake of clarity, without a routing policy that explicitly stops selected routes being re-advertised back to their source, selected routes (e.g. R(A,B)) will also be advertised back to router A.

This part of the history of Route “R” is satisfactory. Route convergence is rapid and without any obvious problems. Timing problems might allow a router (e.g. Router D) to get a route (e.g. from Router B) before the best route (via Router A) arrives. In this case, it will select first R(A,B) and then R(A). This should not. however, cause a problem.

At some time later, Route “R” is withdrawn, and the effect of this is illustrated in Figure 2.

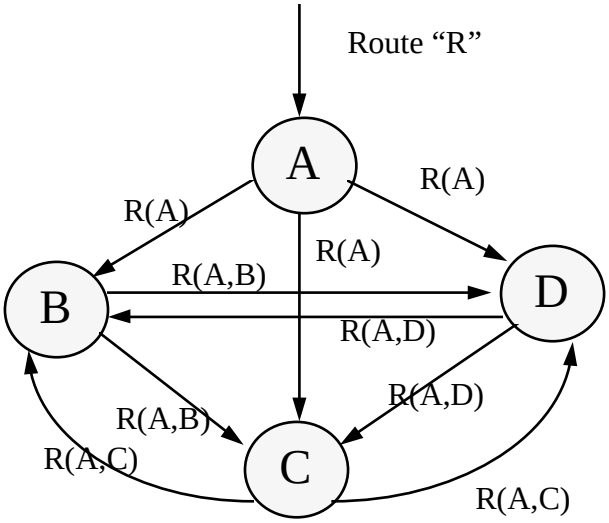

Figure 1 Distribution of Route "R"

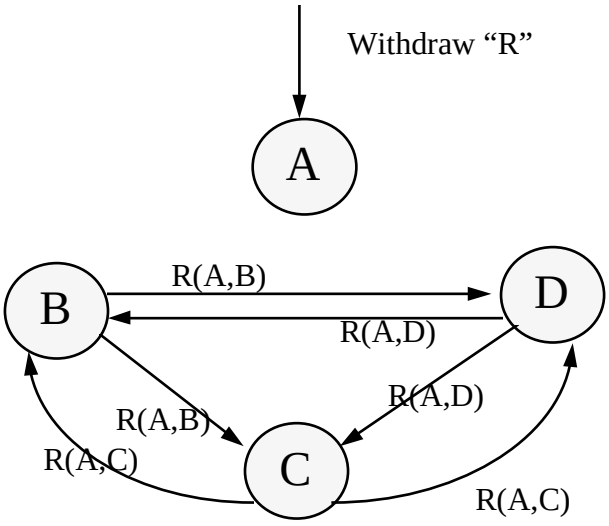

Figure 2 Withdrawal of Route "R"

Router A responds to the withdrawal of Route “R” by withdrawing Route R(A). The withdrawal of R(A) is, for example, received by Router B, which searches its RIB for an alternative route. Two candidate routes exist - R(A,C) and R(A,D). It is assumed that R(A,C) is selected.

Route R(A,C) may be queued for advertisement to Router D and will be advertised once **minRouteAdvertisementInterval** expires. Up and until this time, route R(A,B) still appears available to Router D.

Route R(A,C) may also be queued for advertisement back to Router C if not routing policy is in place to disallow this. If this is true, then for **minRouteAdvertisementInterval** Route R(A,B) also appears to be available to Router C as well. On the other hand, if a rule was in place that prevented the advertisement of R(A,C) back to Router C then, as there is only ever one selected route at any one time, there would be no route "R" to advertise from Router B to Router C, and Route R(A,B) would have to be withdrawn from Router C.

There are thus two possible cases to consider:

(a) when a router does advertise a selected route back to its source, and

(b) when a routing policy rule exists to prevent this.

In case (a), the figure 2 set of false routes does appear to be stable for at least **minRouteAdvertisementInterval**. This is because regardless of which of the two alternative routes, routers B,C and D choose when route R(A) is withdrawn, they will not inform their neighbour BISs until **minRouteAdvertisementInterval** expires.

In case (b), Route R(A,C) will only continue to be available to Router B, if Router C selects Route R(A,D) as its preferred alternative. Similarly, Router D must select R(A,B) as its preferred alternative if it is to continue to advertise Route R(A,D) to Router C. This set of choices is not possible, as clause 7.16.2.1 of ISO 10747 mandates consistent tie breaking rules on BISs which prevents this "vicious circle" of selections. Rapid convergence on no Route to "R" can therefore be expected.

Figure 3 finally illustrates what happens in case (a), when **minRouteAdvertisementInterval** expires. Router B, now advertises R(A,C) to both Routers C and D. Assuming consistent tie breaking rules and that the preference order is R(A,B), R(A,C) and R(A,D), the routes are advertised as in Figure 3. However, Router B, for example, has now lost Route R(A,C), it must ignore Route R(A,B), and hence it immediately withdraws R(A,C). Routers C and D similarly withdraw R(A,B) and a rapid convergence to no Route R occurs.

The comparatively long lived set of false routes illustrated in Figure 2 is certainly undesirable. It will also lead to routing black holes if such a group of routers maintains a false route which other routers continue to select, even if they have an alternative and still valid route to the same destination. This should therefore be regarded as a defect.

2.3.2 Proposed Resolution

ISO 10747 should include requirements that make it a protocol error to advertise a route to another RD which either:

a) Contains the RD's RDI in its path, or

b) Contains the RDI of a confederation which is being entered when the route is advertised to the other RD.

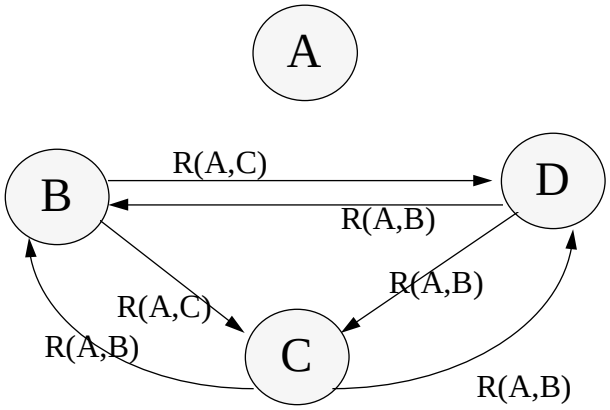

Figure 3 At minRouteAdvertisementInternal

2.4 Applicability of MinRouteAdvertisementInterval

2.4.1 Problem Summary

This particular problem was found when investigating the use of Security Related Information contained in the IDRP Security Path Information.

For example, changes in the security related information of a route are not immediately notified to the adjacent routers to which the route has already been advertised. This is because IDRP deliberately seeks to maintain a stable routing topology at the expense of occasional sub-optimal routing, by the specification of a “Hold Down Timer” (i.e. the **minRouteAdvertisementInterval**) that limits the rate of route re-advertisement. Only new routes and withdrawn routes are not subject to this timer.

However, while this approach is acceptable for routing based on QoS, where “best efforts” is all that is guaranteed, it is not acceptable for satisfying strong policy requirements. Indeed, a defect appears to exist in the ISO standard with respect to the security related information and the Hold Down Timer. For example, the security related information may also be used to convey information on the protection offered by a route. If this protection is reduced, then the route’s users should know about this immediately, so that they may choose an alternative route (if available), that does offer the required protection. Otherwise, data will be discarded at the point at which the protection level falls below that which the user requires.

The problem seems to be that the **minRouteAdvertisementInterval** has not been clearly defined. Its main purpose is to limit the rate of positive changes in routing information in order to avoid unstable routing topologies, but should not limit the propagation of negative changes (e.g. route loss or loss of protection).

2.4.2 Proposed Resolution

Clause 7.17.3.1 should be expanded to include a proper list of exceptions when the **minRouteAdvertisementInterval** does not apply. This list appears to comprise:

a) Routes received from BISs in the same Routing Domain

b) The explicit withdrawal of unfeasible routes

c) Routes with a Security Path Attribute where under the applicable Security Policy, a change in the Security Information contained in the path attribute implies a lowering of protection or otherwise no longer meeting a strong policy requirement.

d) Routes with a reduced Quality of Service.

3. Recommendations

The Draft ATN Internet SARPs already take into account the proposed resolution to 2.4 above. In order to avoid interworking problems between ATN Routers and long lived routing black holes it is recommended that the following changes are made to the draft ATN Internet SARPs, so as to compensate for these other known deficiencies in IDRP:

1 In section 8.3.1.10, the following paragraph and note should be added to avoid ambiguity over the implementation of type 1 authentication:

The type 1 authentication code shall be generated according to the MD4 specification published in RFC 1320.

Note. The interpretation of MD4 given in annex B of ISO 10747 is open to ambiguous interpretation and should not be referenced by implementors.

2 The new section 8.3.1.11 should be inserted, as follows:

8.3.1.11 Restrictions on Route Advertisement

A route shall not be advertised to a BIS in another RD where:

a) The route contains the RD's RDI in its path, or

b) The route contains the RDI of a confederation which is being entered when the route is advertised to the other RD.

Note. This is essential to avoid long lived black holes following the explicit withdrawal of an unfeasible route and when many alternate paths are available (e.g. within an ATN Island Backbone).