WP429

AERONAUTICAL TELECOMMUNICATIONS NETWORK PANEL

ATN Internet Working Group (WG2)

13th Meeting

27th - 30th October 1997

Redondo Beach, USA

Further Guidance Material on the Scaleability of the ATN Routing Architecture

(Presented By Brian Cardwell)

Summary

At the Langen WG2 meeting it was determined that additional GM on the scalability of the ATN Routing Architecture should be included in the ICS Package-1 GM. This material should be based on the information produced in WPs 367, 374 and 391.

This WP reviews the existing GM on the scalability of the ATN Routing Architecture and proposes additional GM based on the WPs listed above.

Introduction.

General

At the Langen ATNP WG2 meeting it was agreed that further ICS Guidance Material should be developed on the subject of the scalability of the ATN Routing Architecture. Three working papers containing relevant material were identified (WP367, 374 &391), two of which were the results of simulations of the routing architectures and the other presenting introductory material on the IDRP Route Server concept.

This working paper will consider the existing guidance material related to the scalability of the ATN Routing Architecture and the content of the three working papers. A proposal for an update to the existing GM will be made for consideration by the Working Group.

Scalability

The scalability of the routing architecture refers to the ability to continue increasing the size of the ATN Internetwork without causing a punitive effect on the original elements of the Internetwork. Thus changes in aircraft traffic or ATN Internetwork topology in one area do not impact upon the network traffic of another. This specifically refers to the containment of IDRP Update traffic; clearly as more Routing Domains and BISs are added to the network there is the potential that network is inundated with route updates and does not achieve routing stability.

Review of material available.

Existing ICS GM, Chapter 5, ATN Routing.

- Section 5.4.4 on Route Aggregation introduces Route Aggregation and Merging and their importance in the implementation of a scaleable Internetwork.

- Section 5.4.5 on Route Information Reduction introduces this function and its contribution to a scaleable Internetwork.

- Section 5.4.6 on Routing Domain Confederations explains the use of RDCs and their roles as containment boundaries for routing information.

- Section 5.9 on Route Selection, Aggregation and Information Reduction gives more detail on the role of these functions and also includes a section (5.9.4) on Scaleable Architectures. The paragraph immediately preceding Figure 5-9 demonstrates how the addition of a Routing Domain has no impact on other routing domains beyond the service provider's.

- Section 5.9.5 on Containment Boundaries and Routing Domain Confederations develops this further by including the RDC to the discussion.

- Section 5.11 on Support for Mobile Systems (20 pages) discusses the impact of mobile RDs on route distribution and how the division of the Internetwork into Islands (RDCs), Island backbones & Home domains, etc. produce a scaleable routing architecture. Section 5.11.6 on the Impact of Routing Updates covers the generation of IDRP Updates, the importance of minimising them and route convergence/stability. This is where the results of the two simulation WPs are relevant.

- Note the final para of 5.11.6.2 refers to the results of a simulation and should be updated to reflect the latest results of that simulation.

ATNP WG2 WP367, "Design of a Scaleable FAA ATN Internet Architecture"

This WP presented the output of a simulation study. From the premise that it is important to minimise the amount of network update traffic to release bandwidth for application data, the study investigated the scalability of an architecture which divided the US into four ATN Islands each containing many a/g routers and a Backbone RD. The Backbone RDs were linked to form single Island backbone.

The key points raised in this WP were:

- IDRP traffic due to aircraft mobility is a linear function of the number of aircraft.
- [Mobile traffic changing A/G BIS within an Island] causes an IDRP notification to be sent to the Backbone only if [an Island's] border is crossed.
- Increasing the number of A/G Routers within an [Island] does not increase routing traffic in other [Islands].
- Increasing the number of A/G Routers within an [Island] does not increase routing traffic at the backbone

The conclusion was the proposed hierarchical architecture was indeed a scaleable architecture.

ATNP WG2 WP374, "IDRP Large Scale Simulation Study: Summary of Results"

This simulation study looked at the relationship of route stability to network topology and sought to determine the optimum topology for minimum time for route convergence. This WP was an update to results from the study that had been presented to WG2 previously (which had led to the generation of some GM).

The earlier study results looked at single ATN Island for the ECAC region and concluded that: the Backbone of the ATN Island should consist of 1 - 3 fully meshed routers to optimise convergence delay, route unavailability period and route update rate; the minRouteAdvertisementInterval time is the only IDRP timer that has a significant impact on route convergence and route update rate.

The follow-up study looked at having multiple ATN Islands and the impact of different interconnection topologies. This study produced the significant conclusion that decomposing the ground part of the ATN into multiple islands results in a scaleable architecture. It also concluded that optimal routing performance was obtained when there was a single adjacency between each pair of ATN Islands.

ATNP WG2 WP391, "The IDRP Route Server Alternative to a Full Mesh Routing"

This WP introduces the IDRP Route Server concept as a solution to the problems of large volumes of route updates when many BISs are interconnected on a WAN and the large number of peering relationships that each BIS would need to maintain. This is all new material.

Conclusions from GM and WP review.

1. The existing GM fully explains the use of IDRP functions IDRP Update, Route Aggregation, Route Merging, Route Information Reduction and RDCs and ATN Islands role in the implementation of a scaleable ATN architecture. This information is presented in a theoretical manner with generic examples.

2. There is no need to add new information on these subjects although the final paragraph of 5.11.6.2 needs to be updated as a result of WP374, specifically the timer value and the topology recommendation.

3. 5.11.6.2 does refer to the implementation of smaller Islands and their interconnection topologies. It may be beneficial to add text an a diagram to illustrate an optimal topology and comment that it is supported by two independent studies.

4. The IDRP route server is not mentioned anywhere in the GM (nor SARPs), however it wasn't necessary to do so as it is an integral part of the IDRP standard. Following the introduction of a diagram showing an optimal topology, it would be worthwhile to briefly explain the use of a Router Server where it is difficult to reduce the number of BIS on a Backbone to an small number.

Recommendation.

1. It is recommended that WG2 review the proposed changes in Appendix A before their inclusion in the CAMAL.

WP429 --- Appendix A - Changes to ICS GM

5.11.6 The Impact of Routing Updates

As indicated in the previous section, a scaleable routing architecture can be developed in support of mobile routing. It is now necessary to consider the factors that limit the number of routes to aircraft that an ATN Router can handle.

Each route known to a router occupies a certain amount of data storage and, while data store can be a limiting factor on the total number of routes handled, it is unlikely to be so in this case. The number of route updates that a router can handle is more than likely to be the limiting factor.

In the ground environment, route updates will usually only occur when changes occur in the local region of the Internet (changes further away are hidden by route aggregation). Typically the introduction of a new Routing Domain or interconnection, or the removal or loss of one of these will cause a change. However, the frequency of update is unlikely to be high.

However, with mobiles, such as aircraft, the situation is very different. Aircraft are constantly on the move, changing their point of attachment to the ATN, and hence generating routing updates. The impact of these updates needs to be minimised if the number of aircraft that can be handled by an ATN Island is to be maximised, and an important and useful feature of IDRP can be exploited in order to help meet this objective.

The KeepAlive timer is used within IDRP to determine the health of a link. This directly controls the frequency which IDRP KeepAlive PDUs are sent on BIS-BIS connections. There is a trade-off concerning the setting of this timer. A small value of this timer will more accurately determine a change in link status, however this will increase the protocol overhead of an already bandwidth limited air/ground resource. The setting of this timer to a small value will also increase the financial cost of the resource. A large value of the Keepalive timer will be less responsive to determine a change in link status, however this will decrease the protocol overhead across the air/ground resource. The setting of this timer to a large value will also decrease the financial cost of the resource. It is recommended that this value be based on operational experience between the various States and Organizations.

5.11.6.1 "Hold Down" Timer Use

Vector distant routing protocols, such as IDRP, typically implement a "hold down" timer, which introduces a minimum delay between the receipt of a route and its re-advertisement. This timer is used to avoid instability due to frequent route changes, and the actual value of the timer is then usually a trade-off between a short timeout to give rapid response and a long timer to keep down routing overhead and minimise instability.

However, under IDRP, routing events that indicate a major change (i.e. new route or loss of a route) are not subject to a hold down timer, only those that report a minor change to an existing route are subject to a hold down timer. This means that IDRP is very responsive to connectivity changes while avoiding instability due to minor changes. For example, consider a simple extension to the previous example, illustrated in Figure 0-1.

In this example, RD4 provides a route to the aircraft, to RD5. When the aircraft loses contact with RD3, RD4 is immediately informed, as there is an effective zero length hold down timer for withdrawn routes. However, while RD4 recognises this event and switches to the route provided by RD2, it does not necessarily inform RD5 of this now minor change to the route immediately (the route still exists, only the detail of the path is different), and anyway, the update must be sent not less than the period

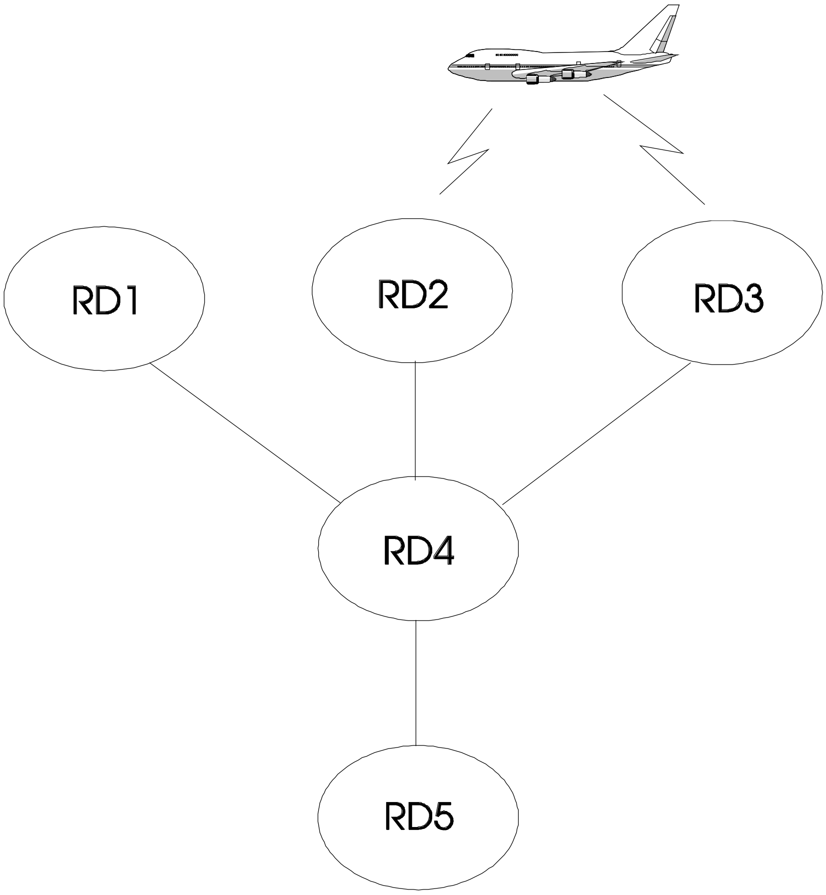

Figure 0-1 Impact of a Hold Down Timer

minRouteAdvertisementInterval since any previous update. In this example, it should be noted that the minor change will not affect RD5's routing decision, as it has no alternatives available.

Sometime later, the aircraft comes into contact with RD1. RD4 is immediately informed as this is a new route. However, even if RD4 switches to this new route, it does not inform RD5 of the change until the **minRouteAdvertisementInterval** has again expired.

This has important implications for the design of an ATN Island. If an Island's air/ground datalinks are all connected to Routing Domains which are themselves adjacent to the Backbone RDC, all connectivity changes will be immediately reported to the Backbone giving a high route update rate. On the other hand, if there are intermediate Routing Domains between the backbone and the Routing Domains connected to air/ground datalinks, then the update frequency can be significantly reduced, without affecting the responsiveness to real connectivity changes.

This is an important benefit derived from using IDRP to support mobile routing compared with, for example, a directory based approach to mobile routing. Under a directory based approach, there would be a central directory server on each ATN Island (c.f. the Backbone), updates on the position of aircraft would be sent direct to the directory, and other routers would consult the directory in order to determine the current location of a specific aircraft. In terms of overhead, this situation is analogous to an ATN Backbone Routing Domain directly connected to each Island Routing Domain with air/ground datalink capability, and the directory has to be able to take the full update rate. IDRP can, however, distribute the update load throughout the ATN Island.

Routes advertised to an aircraft's "Home" are also affected by the hold down timer and, in this case, RDCs and the Hold Timer work together to keep the routing overhead to an absolute minimum.

As an ATN Island is an RDC, routes advertised to other Islands have their path information for the transit through the RDC replaced by a single RDC identifier, and therefore, in many cases, changes in the route will not even be visible to another ATN Island. When changes are visible (e.g. a change in hop count or QoS metric), and such changes can be kept to a minimum by careful network design, then the Hold Timer limits the rate at which such changes can be advertised and prevents minor changes which are also short lived, being exported outside of the Island.

Results from simulation work have shown that the "ideal" setting for the **minRouteAdvertisementInternval** is approximately under one minute (typically 30 seconds). Furthermore, simulation has shown that complex topologies for the ATN Island Backbone should be avoided as they significantly increase the convergence time. Typically, an ATN Island Backbone should consist of a small number of routers linked as a chain with ring shaped topologies avoided. Two independent studies have shown that an hierarchical arrangement of ATN Islands, each with a small number of Backbone BISs, both reduces the volume of IDRP Update traffic and promotes a scaleable architecture. Figure 5-x illustrates such an architecture.

Simulations have also shown that the optimal interconection of ATN Islands is a single direct adjacency between each pair of ATN Islands.

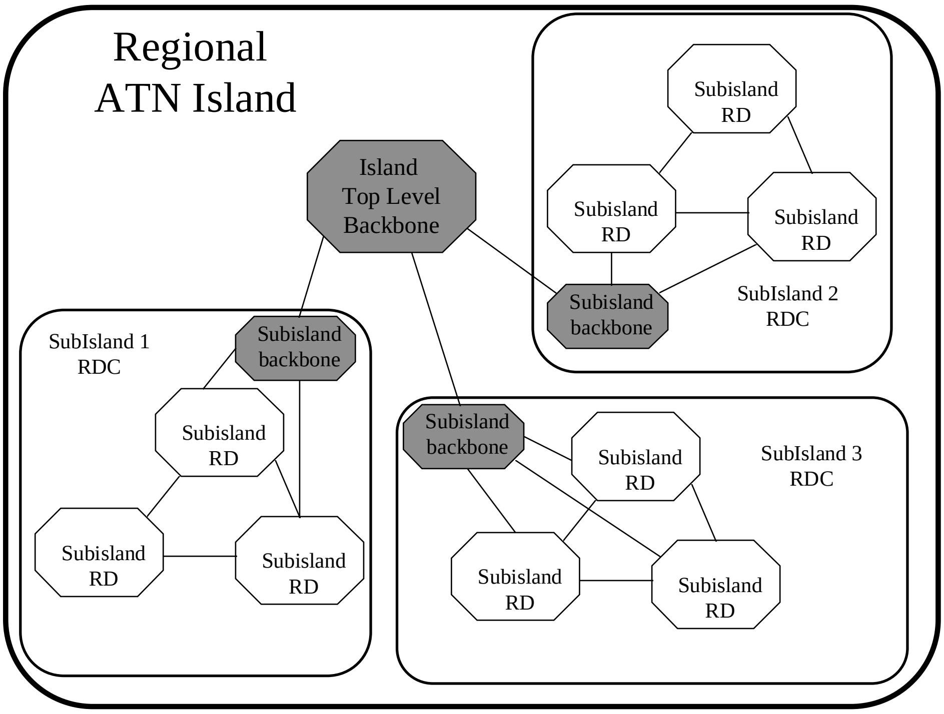

Figure 5-x A Hierarchical arrangement of ATN Islands

Having a small number of BISs on a backbone has been demonstrated to be an optimal arrangement as increasing the number of BISs increases the number of IDRP updates and peer relationships each BIS must handle. However, it may not be possible to produce a topology that satisfies this. Under the circumstances that a Backbone Routing Domain is formed from a number of fully meshed BISs on a WAN, the use of a Route Server may improve the route convergence.

A Route Server is a system that participates in IDRP, but doesn't participate in the actual CLNP packet forwarding. A Route Server is a BIS dedicated to the processing of routes: it acquires routing information from all the BIS connected to a common WAN, performs decision process over this information, and then redistributes the results to the routers. Figure 5.y illustrates the use of a Route Server in a backbone RD.

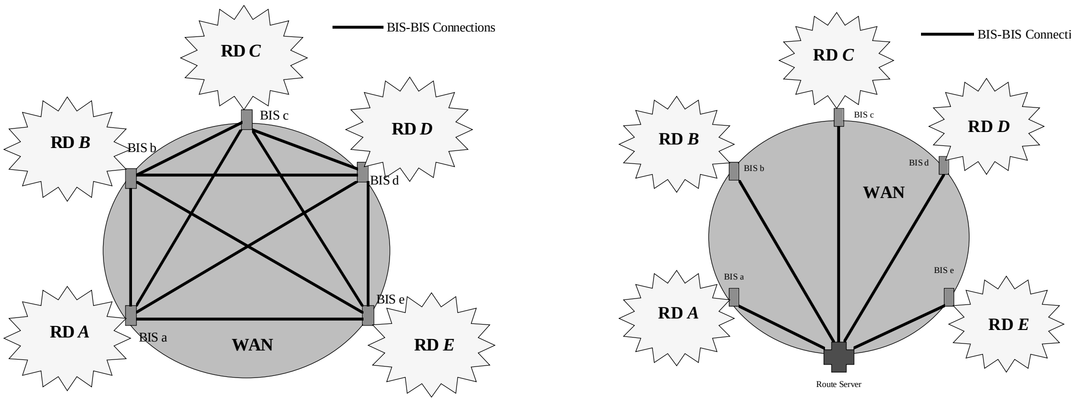

Figure 5-y The use of a Route Server in a fully meshed backbone

The Route Server approach relies on the use of an optional feature of the ISO10747 (IDRP) standard:

- The BISs which are client to the route server would need to support "the generation of the NEXT_HOP attribute in support of route servers" options
- The Router Server is a BIS which does support the "propagation of the NEXT_HOP attribute in support of route servers"
- The support on receipt of the NEXT_HOP attribute is a mandatory IDRP function and is therefore assumed to be supported by every standard/ATN IDRP implementation.

The Route Server approach relies therefore on standard mechanisms and can be used in the ATN provided that the options mentioned above are implemented.