EUROCONTROL

ATNP/WG2/
WP199/
15 December 1995

AERONAUTICAL TELECOMMUNICATIONS NETWORK PANEL

WORKING GROUP TWO

Brisbane 05.02.96-09.02.96

IDRP Route Aggregation and ATN Requirements for Route Aggregation

Presented By Henk Hof

Prepared by Tony Whyman

SUMMARY

Even though Route Aggregation plays a central role in the ATN Internet SARPs, there is still confusion about what it means and limited implementation of Route Aggregation functions - both as a consequence of this confusion and a concern over the cost of such functions. This paper has been prepared in order to remove this confusion and to discuss how Route Aggregation should be implemented and its relationship to routing policy.

DOCUMENT CONTROL LOG

SECTION	DATE	REV. NO.	REASON FOR CHANGE OR REFERENCE TO CHANGE
	15-Dec-95	Issue 2.0	

TABLE OF CONTENTS

1. Introduction

1.1 Background

In IDRP, Route Aggregation, together with Route Information Reduction and Routing Domain Confederations, provides for a scaleable routing architecture that enables the development of worldwide Internets. In the ATN, Route Aggregation is also essential for ensuring that the amount of routing information distributed over low bandwidth air-ground data links is both deterministic and minimal. Route Aggregation mechanisms are also used for the support of air/ground routing in the ground ATN Environment.

However, even though Route Aggregation plays a central role in the ATN Internet SARPs, there is still confusion about what it means and limited implementation of Route Aggregation functions - both as a consequence of this confusion and a concern over the cost of such functions. This paper has been prepared in order to remove this confusion and to discuss how Route Aggregation should be implemented and its relationship to routing policy.

1.2 Scope

This paper discusses the Route Aggregation and Route Information Reduction IDRP mechanisms and how they work with Routing Domain Confederations and Routing Policy. The paper concludes with a discussion on how Route Aggregation relates to the ATN, with implementation considerations.

2. Summary

This paper makes firm recommendations on how Route Aggregation and Route Information Reduction should be implemented by SARPs compliant Routers, and hence proposes additional SARPs requirements in order to clarify this area. It should be noted that these are not new requirements, but are implicit in the existing draft SARPs.

Firstly, it should be noted that automatic Route Aggregation is already required by the draft SARPs for every ATN Ground Router. This is described as Route Merging and is essential to support the policy based routing strategy developed during 1995, in support of multiple air-ground subnetworks. Route Merging is already fully described in the draft SARPs and this paper does not attempt to add any additional requirements in this area.

The recommendations made in this paper apply to Air/Ground Routers where policy based Route Aggregation and Route Information Reduction is required in order to minimise the number of routes advertised from ground to air. Without these functions, a separate route will need to be advertised to each aircraft for every Routing Domain to which connectivity is provided., In practice, this would mean that a route would be advertised to every aircraft for at least every Routing Domain that is adjacent to the Air/Ground Router. Only then would it be possible to ensure that each aircraft has connectivity not only to the current ATC Centre, but also the next one en route. Additional routes would also have to be advertised to support airline requirements. Only by implementing complex policy selection algorithms based on aircraft movements would it be possible to otherwise reduce the number of routes advertised ground to air.

An IDRP route is greater than 100 bytes long. Thus, for example, in the case of Mode S, a scan would be taken up for every route advertised. In a complex ground environment, such as that which exists in Europe today, this would imply that without Route Aggregation in the Air/Ground Router, several minutes could elapse during the Route Initiation phase. Furthermore, the time taken for Route Initiation would not be bounded and would increase linearly with the Air/Ground Router’s connectivity.

There is thus strong justification for implementing Route Aggregation procedures in Air/Ground Routers, as the alternative is either excessive overhead on the air-ground data link or complex policy decisions based on knowledge of aircraft movements.

All ATN Routers must anyway implement the procedures specified in ISO 10747 for the aggregation of IDRP path attributes mandated for ATN use. This is already required for Route Merging and is not a new requirement.

The recommendations made in this paper apply specifically to air-ground routers and derived from the existing requirements in the draft SARPs are:

1. That, when choosing the routes to be advertised to a given Airborne Router, all Air/Ground Routers support the policy based selection of routes. The selection of candidate routes for aggregation should be according to a filter on each route's destination, with a combination of inclusion and exclusion filters e.g. to select all routes to NSAP Address Prefixes within the local ATN Island, but excluding those to the local Administrative Domain.

2. That all Air/Ground Routers support policy based Route Information Reduction, and Routing Policy rules are required to specify when a set of NSAP Address Prefixes is replaced by a shorter NSAP Address Prefix. Two types of rules are recommended:

   a) The explicit replacement of a set of NSAP Address Prefixes by another shorter NSAP Address Prefix, only when all members of the set are present; or.

   b) The explicit replacement of a set of NSAP Address Prefixes by another shorter NSAP Address Prefix when any members of the set are present.

Other ATN Routers, for example, at the boundaries of ATN Islands, are likely to require policy based Route Aggregation and Route Information Reduction functions, in order to allow the ATN Ground Environment to develop in a scaleable fashion. However, this is not a near term requirement.

The remainder of this paper is concerned with discussing the role of Route Aggregation and Route Information Reduction, and the support given to them by Routing Domain Confederations (RDCs), and how these concepts are used in the ATN.

3. Route Aggregation and IDRP

3.1 The Scaleability Problem

In principle, a Router has a simple task - it extracts the destination address from the header of each packet, and uses this as an index to a table to determine:

- firstly, on which network to forward the packet; and,

- secondly, the address of the Router or End System on that network, to which the packet is to be forwarded.

The packet may then be forwarded. However, such a table will always have a finite capacity and, with a simple approach that just lists every possible destination in a routing table, the Router with the smallest capacity routing table will then place an effective upper bound on the size of any network.

Moreover, when dynamic adaptive routing protocols are used to maintain routing information, if an entry has to be maintained in a routing table for every possible destination, then as a network grows, the number of routing information updates will grow rapidly and soon become excessive. For example, when someone switches off a PC in Australia, there should really be no need for routers in Europe to be explicitly informed of this.

Some kind of structured approach has thus always been regarded as necessary, hiding information from Routers that do not need to know it and attempting to minimise the information that a Router needs to know.

3.2 Address Assignment

Typically, such a structured approach starts with the hierarchical assignment of network addressing, assigning such addresses in either regional or organisational hierarchies. The obvious example of hierarchical assignment occurs in the telephone network, where individual subscriber numbers are assigned relative to the number assigned to their local telephone exchange, which are in turn assigned relative to area codes and then country codes. A nested set of addressing domains is thereby formed, with each country being an addressing domain, then each telephone area and finally each telephone exchange, also being addressing domains. With this structure, different administrators can assign telephone numbers in different domains without interfering with each other. Calls can be routed first by country, then by area, and finally within each telephone exchange, without each level in the routing hierarchy having to be aware of the topology of other levels.

In such a structured approach, all network addresses assigned by a given network administrator have a common prefix, that being the prefix allocated to the network administrator for address allocation within their domain. All such addresses assigned by a network administrator form an addressing domain.

The routing information exchange protocols need to recognise such an addressing hierarchy, if they are to take advantage of it, and to distribute routing information such that:

a) between addressing domains, routing information is limited to describing routes to each addressing domain, and characterising each addressing domain by the address prefix common to all addresses within the domain; and,

b) within addressing domains, routing information may then be more detailed, describing routes to each assigned address (which may themselves be prefixes assigned to subordinate addressing domains).

Essentially a containment boundary needs to be defined, that contains the distribution of detailed routes within the scope of an addressing domain. In practice, addressing domains relate to organisational hierarchies, while the containment boundaries recognised by routing protocols take into account network topology. The containment boundaries recognised by routing information exchange protocols may not therefore always coincide exactly with addressing domain boundaries, although, more often than not, they do.

3.3 Routing Domains

In OSI routing, the most important containment boundary is formed by the Routing Domain. A Routing Domain is formally a grouping of End Systems and Intermediate Systems under a common network administrator. A Routing Domain also coincides with at least one addressing domain, but may also enclose several addressing domains. The systems contained within a Routing Domain will therefore have either a single common address prefix, or at least have an address prefix chosen from a limited set. In either case, the address prefix will be unique to the Routing Domain.

Routing Domains are important because they contain systems that are subject to a common routing policy and are within a common domain of trust. Within a Routing Domain it is possible to use a connectivity based routing protocol such as that defined in ISO 10589, while between Routing Domains a policy based routing protocol, such as that defined in ISO 10747 has to be used.

In a Routing Domain that implements ISO 10589 internally, a two level approach to containment is applied within the Routing Domain, with the Routing Area forming a subsidiary containment boundary. The approach taken by ISO 10589 is not directly relevant to this discussion, but is useful for background understanding, and is hence presented as an appendix to this paper.

3.4 Inter-Domain Routing

The Inter-Domain Routing Environment may be assumed to comprise multiply interconnected Routing Domains. Some of these are Transit Routing Domains (TRDs), which will relay packets between other Routing Domains, whilst others are End Routing Domains (ERDs) which never relay between other Routing Domains, but may still have connections with many other RDs.

These RDs will be organised into many, possibly overlapping hierarchies, representing regional and organisational groupings. Service providers will operate TRDs, which will support many customers, perhaps having a single ERD, or perhaps their own grouping of TRDs and ERDs. Service Providers will themselves be interconnected, and many users will interconnect with multiple service providers.

The ISO Inter-Domain Routing Protocol (IDRP) is designed to operate in this environment. Recognising both the need to support multiple overlapping hierarchies and the lack of any common domain of trust, IDRP deals in routes, unlike the connectivity information distributed by ISO 10589. An IDRP Route comprises a set of destinations and information about the path that the routes takes through one or more RDs. The destinations are NSAP Address Prefixes, and a Router implementing IDRP aims to forward packets along the route which contains in its set of destinations, the NSAP Address Prefix that provides the longest match with the packet’s destination NSAP Address. The destination of a route to an Routing Domain is the unique address prefix(es) that is common to all addresses within the Routing Domain.

IDRP routes are advertised from one Routing Domain to another, not just because the connectivity exists, but because organisational policy permits the route to be advertised, makes available the path offered by the route, and the resources that packets following this route will consume. Similarly, a Router is not required to always accept a route it receives - organisational policy may be to reject it, and organisational policy is also used to select between routes to the same or overlapping sets of destinations.

IDRP also gives Routing Domains a name. This is the Routing Domain Identifier (RDI), and is totally separate from the unique address prefix(es) common to all systems within the Routing Domain. An RDI is a Network Entity Title (NET) and is therefore syntactically a single address. It will typically be assigned relative to the Routing Domain’s common NSAP address prefix, but does not have to be. It is used by IDRP in the RD_Path attribute that forms part of each route’s path information, and records that the route passes through the Routing Domain. It is used for loop detection purposes and may also be referenced by routing policy.

3.5 Containment beyond the Routing Domain

However, while such features are necessary to support routing in the inter-domain environment, they do not of themselves aid scaleability. For example, if a Routing Domain generates a route to internal destinations (i.e. with a destination that is the unique NSAP Address Prefix common to all NSAP Addresses in the Routing Domain), and then that route is relayed to other RDs, with routing policies designed to ensure maximum connectivity, every Router at each Routing Domain's Boundary (i.e. a Boundary Router) will eventually be informed about that route, and will make a corresponding entry in their routing tables.

If every Routing Domain generates a similar route to internal destinations then, in consequence, every Boundary Router will have a routing table entry for every other Routing Domain. RD's at the periphery of an internetwork may be able to implement local policies that exclude routes to destinations for which they have no interest, but Service Providers are not able to implement such policies. An upper limit on the number of RDs in the Internet will be the number of routing table entries that the smallest Service Provider can support.

To avoid this problem, the same sort of techniques are required as are applied by ISO 10589. Essentially, containment boundaries, linked to the hierarchical allocation of NSAP Addresses, are necessary to provide limits on the scope of routing information distribution, without preventing communication between the systems in different areas of containment. Then only reduced routing information need be exchanged between routers within different containment boundaries.

In IDRP, the Routing Domain Confederation provides a generalised containment boundary that reflects the nature of the inter-domain environment. This is complemented by Route Aggregation and Route Information Reduction procedures that, respectively, permit the merging of routes, in order to minimise the number of routes, and, the merging of NSAP Address Prefixes in route destinations, in order to minimise the information distributed to each router. Route Aggregation and Route Information Reduction provide the mechanisms by which routing information is reduced prior to it being advertised outside of an RDC.

3.6 Routing Domain Confederations

A Routing Domain Confederation (RDC) is a set of Routing Domains. It is a very general concept and RDCs may be nested within one another, and may also overlap. RDCs only have one functional role, that is to contain the expansion of trace information in a route's path. Otherwise, they provide containment boundaries that can be readily referenced by routing policy.

Like Routing Domains, RDCs also have names, and again these are NETs, and are also called RDIs. The RDI of an RDC is also used in the IDRP RD_Path to record when a route enters an RDC and, the RDCs that the route has passed through.

The use of RDCs in controlling trace information is an important contribution to scaleability. Every time a route is advertised by a Routing Domain, the Routing Domain's own unique Routing Domain Identifier (RDI) is added to the route's trace information held in the RD_Path path attribute. This attribute keeps track of which RDs the route has passed through, in order to prevent routing loops. However, when a route exits an RDC, the RDIs of all RDs that are within the exited RDC are removed and replaced by the RDI of the RDC itself. This has the dual role of both reducing the overhead of the trace information, and of preventing the route from ever re-entering the RDC, as to do so would constitute a routing loop.

RDCs thus by themselves have an important role to play in reducing the overhead of distributing routes in large internets. However, by ensuring that once a route has left an RDC it cannot re-enter it, they also provide a useful basis for routing policy in scaleable internetworks, as Route Information Reduction can then be implemented without risk of ambiguous routing.

3.7 Route Aggregation

Route Aggregation is formally the process by which two or more routes are merged together to form a single route. Each IDRP path attribute has route aggregation rules associated with it, determining the mechanics of route aggregation. Route Aggregation is applied when:

a) Local Routing Policy selects a groups of routes for aggregation prior to their advertisement to another Routing Domain; or,

b) Two or more routes selected for advertisement to an adjacent Routing Domain have an identical set of NSAP Address prefixes as their destination, have the same distinguishing path attributes, including the security path attribute, but have different security information in their paths.

The former case is generally used in conjunction with Route Information Reduction in order to reduce the number of routes being advertised, while the latter is necessary to ensure proper routing when routing control procedures are in effect. This is because it is not possible to advertise two route to an adjacent Router with the same destinations and distinguishing path attributes, without one being assumed to be a replacement to the other, rather than as a different route.

3.8 Route Information Reduction

Route Information Reduction is the process by which the set of NSAP Address Prefixes contained in the destination of a route are replaced by a smaller set of shorter NSAP Address Prefixes. This may occur because:

a) The set of NSAP Address Prefixes is mathematically complete, and the shorter NSAP Address Prefixes provide equivalent routing information; or,

b) Local Routing Policy specifies that a given set of NSAP Address Prefixes (possibly a subset of those present) are replace by a shorter NSAP Address Prefix; or,

c) Local Routing Policy specifies that when *any* members of an identified set of NSAP Address Prefixes are present in a route’s destination, those members of the set that are present are replaced by a shorter NSAP Address Prefix.

Route Information Reduction usually, but not necessarily, occurs after Route Aggregation and ensures that detailed routing information to groups of Routing Domains, is hidden from outside the group. In conjunction with intelligent allocation of addresses, it thus plays the crucial role of ensuring that the size of an Internet is not limited by the capacity of the smallest Service Provider. For example, if Route Information Reduction takes place at the boundaries of an RDC that also encloses an Addressing Domain from which RDs within the RDC are assigned NSAP Address Prefixes, then the number of RDs within a given RDC can increase without affecting Routers outside of the RDC.

However, Route Information Reduction is not tied to RDC boundaries. In the first case listed above, the information reduction clearly can take place anywhere without ambiguity. The second case, is essentially similar. A Network Administrator defines such a rule when only a limited number of NSAP Address Prefixes have been allocated relative to another shorter prefix. It is essentially saying that “when you have all these prefixes in the destination of a route, then you can replace them by this shorter prefix without ambiguity, because no other prefixes relative to this shorter prefix can exist”.

It is, though, the final case that is the really important one for scaleability, and the one that really should be used in conduction with RDCs.

3.9 Routing Policy and Route Aggregation

A Routing Domain provides a containment boundary and, regardless of which systems are switched on within the Routing Domain, there is no need to qualify the route to the Routing Domain, depending on which systems are active. The route advertised by the Routing Domain is a route to the unique NSAP Address Prefix(es) assigned it. This is always an unambiguous route.

RDCs can also be used as containment boundaries, and in a recursive fashion. They therefore enable scaleability of the inter-domain environment.

Firstly, when a group of Routing Domains share a unique NSAP Address Prefix, they should form the membership of an RDC. This RDC is then a containment boundary that coincides with the addressing domain from which those NSAP Address Prefixes were assigned.

Secondly, at the boundaries of the RDC, the routing policies should be specified such that all routes that originate within the RDC are first aggregated, and then Route Information Reduction applied. The Route Information Reduction should be of type 3.8(c) above, and set up so that any NSAP Address Prefix(es) for destinations inside the RDC are replaced by the single unique NSAP Address Prefix shared by all systems within the RDC.

When the above is in place, such an RDC appears like a single Routing Domain. The internal detail is hidden from those outside, and the membership of the RDC can grow and change without affecting those outside of it. Even if whole Routing Domains are “switched off”, this fact is hidden to the outside of the RDC, given the form of Route Information Reduction applied.

As RDCs can be nested within each other, this form of information hiding can occur many times over, perhaps hiding such information on first a company, then a regional, and then on a Country basis.

What’s more, the same process can happen in reverse. At the boundaries of an RDC, routes entering the confederation can be aggregated and their destinations replaced by a single NSAP Address Prefix for the “rest of the world” - although if there is more than one entry point to the RDC, such a mechanism may require careful co-ordination if routing to outside of the RDC is not to suffer from unexpected behaviour.

4. Route Aggregation and the ATN

4.1 In the Ground-Ground Environment

The draft ATN Internet SARPs do not mandate any route aggregation in support of ground-ground routing, although clearly this will become necessary as the ATN grows. However, route aggregation is mandated in support of air/ground routing.

4.2 Supporting Air/Ground Data Links

The draft SARPs defines a number of RDCs in support of air-ground routing. These are containment boundaries, and, together with specific use of Route Aggregation and Route Information Reduction, support the requirements for air-ground route information distribution.

4.2.1 The Fixed ATN RDC

The draft SARPs define the “Fixed ATN RDC” as comprising all ATN RDs other than mobile RDs and also specify an addressing plan that defines separate NSAP Address Prefixes for ground based and mobile systems. The Fixed ATN RDC is a containment boundary, containing all ground systems. Furthermore, all systems within this RDC have a common NSAP Address prefix. The Fixed ATN RDC has three main purposes:

1. To minimise the trace information sent over the air/ground data link. As this confederation is exited whenever a route is advertised from an air/ground router to an airborne router, all trace information concerning the ground environment is removed and replaced by the single RDI of the Fixed ATN RDC.

2. To ensure that erroneous routes between two air/ground routers and via an airborne router, are prevented. This is a simple consequence of the fact that routes cannot re-enter a confederation.

3. To support the recommendation in sections 3.7.1.3 and 3.7.3.3 of thedraft SARPs to advertise to an airborne router an aggregated route to all ATN Destinations on other ATN Islands. It is appropriate to implement this recommendation when the Air/Ground Router has connectivity, directly or indirectly, with the Global ATN Backbone envisaged in the draft ATN SARPs.

In the last case, an aggregation rule will be required to select routes to destinations outside of the local ATN Island for aggregation into a single route. As the route is now exiting the Fixed ATN RDC it is also possible to apply a Route Information Reduction rule replacing any NSAP Address Prefix(es) for a destination outside of the ATN Island, with the common address prefix for all ground systems.

The effect of this is to offer a single route to “the rest of the ATN”, with a single address prefix in its set of destinations, and a single entry in the trace information. This enables the availability of worldwide connectivity to be efficiently advertised to an airborne router.

4.2.2 The ATN Island RDC

The ATN draft SARPs define the ATN Island RDC as comprising all ATN RDs within the same ATN Island. This is illustrated in Figure 1, which also illustrates the flow of routes within an ATN Island. Like the Fixed ATN RDC, the ATN Island RDC is also a containment boundary, although the draft SARPs do not specify that there is a common NSAP Address for all systems within an ATN Island RDC, nor is there a recommendation to this effect. This is arguably a defect in the current SARPs, as inefficient use of air/ground data links may result if there is no more than one or two unique NSAP Address Prefixes for all the systems within an ATN Island. The ATN Island RDC has two main purposes.

1. To minimise the trace information on routes advertised between ATN Islands.

2. To support the recommendation in sections 3.7.1.3 and 3.7.3.3 to advertise to an airborne router an aggregated route to all systems in the local ATN Island RDC.

In the last case, an aggregation rule will select routes to destinations within the ATN Island, other than those to the local Administrative Domain, and aggregate them together into a single route for advertisement to an airborne router. Provided that all systems within the ATN Island RDC share a unique NSAP Address Prefix, then Route Information Reduction will also be appropriate here. This requires a rule that replaces any NSAP Address Prefix(es) for a destination within the ATN Island, other than those of the local Administrative Domain, with the common address prefix for all ground systems.

The effect of this will be to offer a single route to the "rest of the ATN Island". However, the efficiency of the propagation of this route will be limited if there is no single unique address prefix for all systems within the ATN Island.

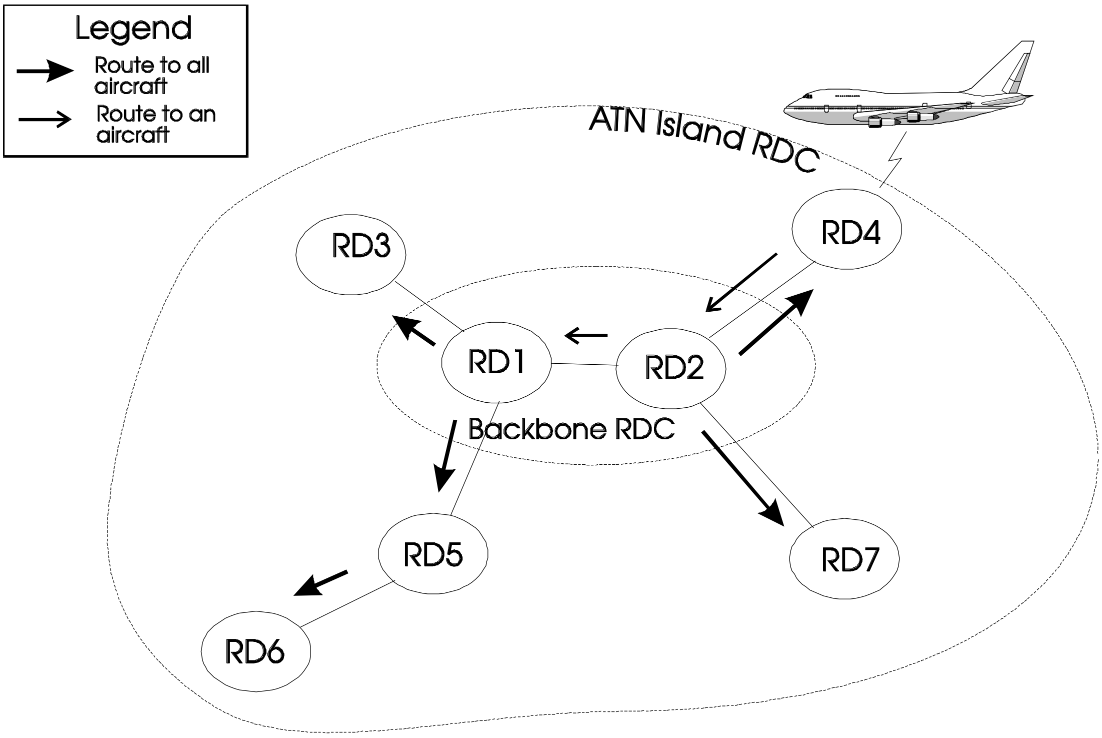

Figure 1 The ATN Island RDC

4.2.3 Administration RDCs

The draft SARPs do not require that administrations and organisations that own/operate multiple RDs, group them together into an RDC. However, such an RDC will be useful in providing a containment boundary for each administration or organisation's systems. The current addressing plan does specify a common NSAP Address Prefix for all of an Administration or Organisation's systems, and such an RDC would therefore be the containment boundary for systems with addresses assigned relative to this NSAP Address Prefix. However, this part of the addressing plan is not fully compatible with the goal of having a common address prefix for an ATN Island and may need to be reviewed.

4.2.4 The ATN Island Backbone RDC

The draft ATN SARPs define the ATN Island Backbone RDC as comprising the backbone RDs within an ATN Island. It provides a containment boundary for the distribution of routes to mobiles.

Routers within a Backbone RDC are required to implement a routing policy that does not permit the advertisement of routes to individual mobile RDs, to be advertised outside of the confederation, except to the preferred route to an aircraft's Home. Instead, such routers advertise a default route to all aircraft.

This RDC provides a useful boundary that may be referenced by routing policies, but no more. This is not intended to be an Addressing Domain boundary. Furthermore, the Default Route to all mobiles that is advertised from the Backbone Routing Domain to the rest of the ATN Island is not formed by Route Aggregation, but is instead generated *a priori*.

This is because, such an aggregated route would contain RDIs for (possibly all) RDs in the ATN Island. Route loop prevention functions would therefore prevent its advertisement to any other Routing Domain in the ATN Island, therefore defeating its purpose. The default route to all mobiles is therefore totally separate from individual routes to mobile systems, and is not an aggregated route. For similar reasons, the default route generated by “Home” RDs, also not formed by Route Aggregation.

4.2.5 Route Merging

Route Merging is already specified by the draft SARPs and is a specific example of Route Aggregation. Route Merging takes place automatically (i.e. without reference to any routing policy rules), and is necessary to support routing via different air/ground subnetwork technologies on a per application basis. Without Route Merging, an intermediate router that had a choice of routes to the same aircraft via different air-ground subnetworks, would have no choice but to discard one of those routes. In consequence, any “downstream” systems would be denied access to this subnetwork.

Route Merging is also specified with simplified rules for aggregation of the RD_Path path attribute. This simplification was introduced as a short term measure to counter concerns over whether the full algorithm could be validated in time for the ATN Panel meeting.

This form of Route Aggregation corresponds to that described in 3.7 (b). There is no requirement for this aggregation to take place at RDC boundaries, neither is it likely to be used with Route Information Reduction.

5. Implementation Considerations

Except for Airborne Routers which, as single routers serving ERDs, never need to be involved in Route Aggregation or Route Information Reduction, all ATN Boundary Routers are required to support Route Aggregation. However, support for Route Information Reduction is less widely needed. Also, there are different levels of Route Aggregation support required. The following is an interpretation what the draft ATN SARPs require:

1. All ATN Inter-Domain Routers other than Airborne Routers must implement the functions for Route Aggregation i.e. those functions that perform the aggregation of two or more routes into a single route, according to the provisions for path attribute aggregation specified in ISO 10747. Such routers must be able to perform Route Aggregation, as required by 3.7 (b). This is known by the draft SARPs as Route Merging.

2. Air/Ground Routers are required to support Route Aggregation according to 3.7 (a) i.e. to aggregate routes selected by routing policy rules. Such rules need to be applied on a per adjacent BIS basis, as aggregation is only required on air/ground data links. However, it could be acceptable to apply such rules depending on which RDCs are being exited. For ATN use, the selection of candidate routes for aggregation will be according to a filter on each route’s destination, with a combination of inclusion and exclusion filters e.g. to select all routes to NSAP Address Prefixes within the local ATN Island, but excluding those to the local Administrative Domain.

3. The ATN Draft SARPs do not require that policy based Route Aggregation is supported elsewhere. However, it should be viewed as desirable whenever a router is sited on a confederation boundary, especially when such confederations are associated with address allocation boundary.

4. Route Information Reduction is required only in Air/Ground Routers; it is not required in support of 3.7 (b) type Route Aggregation, but is required to minimise the information transferred over air-ground data links. Route Information Reduction

is rule based, with routing policy rules required to specify when a set of NSAP Address Prefixes is replaced by a shorter NSAP Address Prefix. In general, two types of rules are required:

a) The explicit replacement of a set of NSAP Address Prefixes by another shorter NSAP Address Prefix, only when all members of the set are present;, or

b) The explicit replacement of a set of NSAP Address Prefixes by another shorter NSAP Address Prefix when any members of the set are present.

5. The ATN Draft SARPs do not require that Route Information Reduction is performed elsewhere. However, it should be viewed as desirable whenever a router is sited on a confederation boundary, especially when such confederations are associated with address allocation boundary.

Table 1 summarises the support requirements for the two forms of Route Aggregation and Route Information Reduction, in ATN Routers.

6. Recommendations

1. It is recommended that the ATN draft SARPs are clarified with respect to Route Aggregation and Route Information Reduction by adding SARPs that make clear the policy rules needed in support of these functions in order to implement the existing provisions of the draft SARPs.

2. It is recommended that a table of the form of Table 1 is also included in the draft SARPs in order to clarify the support requirements for each class of ATN Router.

Class	Name	Route Merging	Policy Based Route Aggregation	Policy Based Route Information Reduction
1.	Static Router			
2.	Level 1 Router			
3.	Level 2 Router			
4.	Ground-Ground Router	✓		
5.	Air/Ground-Router (ground based)	✓	✓	✓
6.	Airborne Router with IDRP	✓		
7.	Airborne Router without IDRP	✓		

Table 1 ATN Routers and Route Aggregation

Appendix A - Scaleability within a Routing Domain and ISO 10589

In ISO 10589, is specified as a routing information exchange protocol for use within a Routing Domain. Under ISO 10589, each network address (specifically an OSI NSAP Address) assigned to End Systems is assumed to be in two parts: the first part - the address prefix - is an Area Address, and, the second part, is a system identifier. End Systems and Routers are then grouped together into Routing Areas, where all the systems in a given Routing Area are assigned NSAP Addresses such that the Area Address part of their NSAP Address is one of a small set of Area Addresses that are uniquely assigned to the Routing Area. As long as the System Identifier part is unique within the Routing Area then each system in a Routing Area can be assigned a unique NSAP Address.

The End Systems and Routers within a Routing Area are interconnected by whatever networking technology is appropriate, and the Routing Area forms a containment boundary. Within the Routing Area, the Routers have to keep track of where each active End System and Router is within the Routing Area, and the ISO 10589 routing information exchange protocol helps them to do this and to keep track of the Routing Area’s topology. The size of an individual Routing Area is therefore always limited by the capacity of its smallest Router’s routing table. However, this is not a limitation on overall network size, because the Routing Area’s Routers only need to know the Area Addresses of other Routing Areas in order to route out of area packets.

Indeed, knowledge of other Area Addresses is further limited to specially configured “Level 2” Routers, at the boundaries of Routing Areas. The remaining “Level 1” Routers within the Routing Area need only to recognise that a packet’s destination contains an Area Address that is out of area, and then forward the packet to the nearest Level 2 Router.

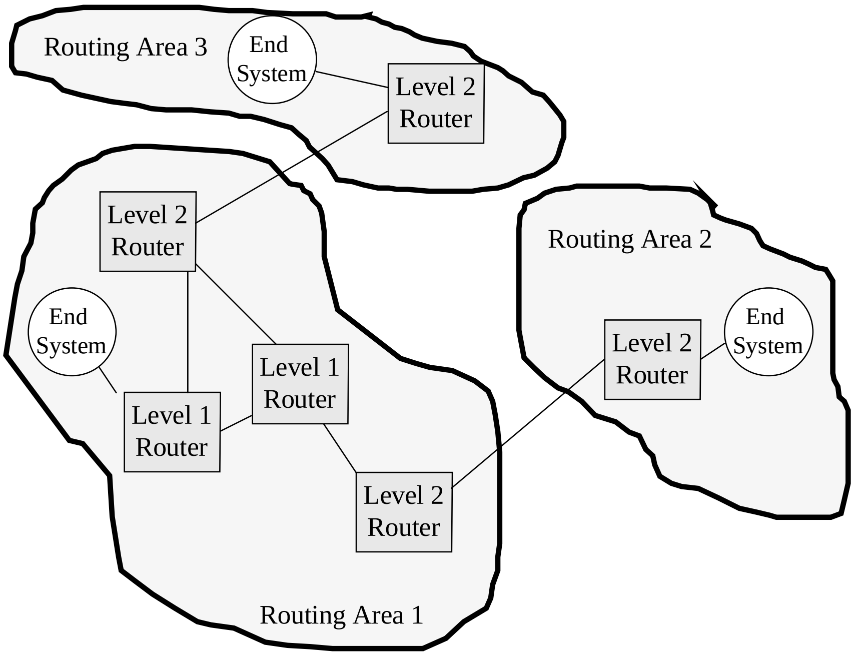

Figure 2 Example Routing Area Interconnection

Routing Areas are typically grouped together within a single organisation to form a Routing Domain, and the Level 2 Routers within a single Routing Domain are interconnected, as the Level 1 Routers are within a Routing Area. This is illustrated in Figure 2. The ISO 10589 protocol permits Level 2 Routers to keep track of the topology of their interconnection - the

Level 2 Domain - and to find out the Area Addresses supported by each Level 2 Router. It is therefore possible for Level 2 Routers to forward “out of area” addressed packets to a Level 2 Router in their destination Routing Area, and thence to the destination system.

As with a Routing Area, the size of a Routing Domain is going to be limited by the capacity of the smallest Level 2 Router, but it is more likely that organisational boundaries are going to be reached first. ISO 10589 assumes that all Routers work within a common domain of trust and are subject to a single administrative policy. In general, it is not suitable for inter-organisational routing, except in certain limited cases.

For example, ISO 10589 does permit certain Level 2 Routers to be configured as Boundary Routers. These Boundary Routers have *a priori* knowledge of the Area Addresses assigned to Routing Areas in other Routing Domains, and interconnections with similar Boundary Routers in those Routing Domains. Ideally, all the Area Addresses assigned to Routing Areas within a single Routing Domain share a single (and shorter) unique NSAP Address Prefix, and then such Boundary Routers only need have *a priori* knowledge of the NSAP Address Prefix for the Routing Domain as a whole, rather than of each Routing Area within the Routing Domain.

The ISO 10589 protocol then permits the Level 2 Boundary Routers to inform other Level 2 Routers of their connectivity with remote Routing Domains, and the NSAP Address Prefixes that uniquely identify all destinations within such Routing Domains. It is therefore possible to forward packets addressed to destinations in other Routing Domains, through the Level 2 Domain and to Boundary Routers that can pass those packets to the Level 2 Domain of the destination Routing Domain.

This is clearly a very limited approach to routing between Routing Domain - i.e. inter-domain routing. This is because it is generally impractical to have direct interconnections with anything more than a small number of adjacent Routing Domains, and anyway, expansion is still limited by the smallest Level 2 Router. This is because when there is more than one Boundary Router in a Routing Domain, each Level 2 Router must explicitly know which Boundary Router supports routes which adjacent Routing Domain. A more structured approach is therefore necessary to inter-domain routing.