AERONAUTICAL TELECOMMUNICATION NETWORK PANEL

Working Group 2

Munich, Germany 24 - 28 June 1996

Proposed Draft Guidance Material For Route Aggregation in the ATN

Presented by HENK J. HOF

Prepared by Tony Whyman

SUMMARY

This Working Paper provides draft Guidance Material on the subject of Route Aggregation. It is a development of the material submitted for comment in IP272.

DOCUMENT CONTROL LOG

SECTION	DATE	REV. NO.	REASON FOR CHANGE OR REFERENCE TO CHANGE
	21-Jun-96	Issue 1.0	

TABLE OF CONTENTS

1. An Informal Introduction to Route Aggregation

1.1 What is Route Aggregation?

Route Aggregation is one of those subjects guaranteed to empty a room. Far from being the kind of subject that can be used to break the ice at parties, it seems to have as much social value as combining religion and politics in the same sentence. However, it is very relevant to the building of big Internetworks because, without Route Aggregation, we have no way of controlling the amount of routing information that our Routers have to deal with. And, if we can't control how much routing information is exchanged then, ultimately more routing information will end up getting exchanged than our Routers can cope with, and that will be the limit on the size of network we can build.

So, Route Aggregation is worth knowing about. But what is it? How does it help us build a big Internet, moreover, why is it relevant to the ATN?

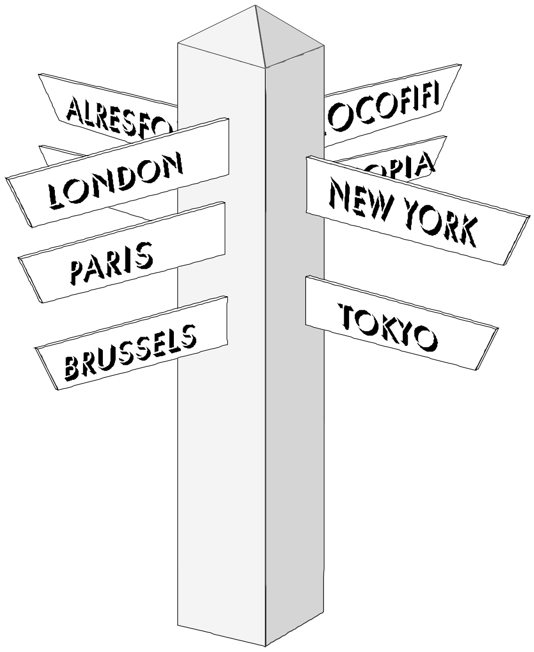

Figure 1-1 Signposting the Way

Well, look at the signpost alongside in Figure 1-1, and imagine being confronted with it at a road junction. If you are going to one of the big cities indicated on it, then you're in luck. It points you in the right direction. But, if you are not, what do you do? Complain to the person that erected it?

Perhaps you do. You want to go to Berlin, and you're the kind of person that complains strongly if things aren't right. The person responsible for the signpost, reacts to customer demand and adds a sign for Berlin. Off you go, a satisfied customer.

The same then happens for people wanting to go to Rome, Toulouse, Sydney, Singapore, Peking, Cape Town, Rio de Janeiro, Seattle, Moscow, Dublin, Brisbane, Winchester, Prague, Bristol, Athens, Anchorage, Stornoway, Oslo, St Petersberg, and so on, until there is no further room on the signpost to hang another sign. What does our poor Signpost Manager do now?

He could just erect a bigger signpost, but if he's bit cleverer, he may just realise that the problem is not one of insufficient signpost real estate, but really it's the granularity of information that is being provided. After all, London, Paris and Brussels are all in Europe, and hence could be replaced with a single sign indicating the direction to Europe, along with all the other cities and towns in Europe that are individually listed on the signpost.

In fact, this is a really bright idea, as it is not just the European cities that can be picked off in this way, but so can the Asian cities, the American ones, the African ones, and so on. Only those that really are local (i.e. on the same continent) need to be explicitly mentioned. What our bright signpost manager has realised is that his customers don't really need detailed information on the route for their individual destinations. There are only a few directions in which they can go anyway and, when he labels each direction with a suitable collective noun or group name, that properly and unambiguously describes what is reachable in that direction, the signpost's users will get all the information they need. After this exercise in information reduction, our signpost ended up much like that in Figure 1-2.

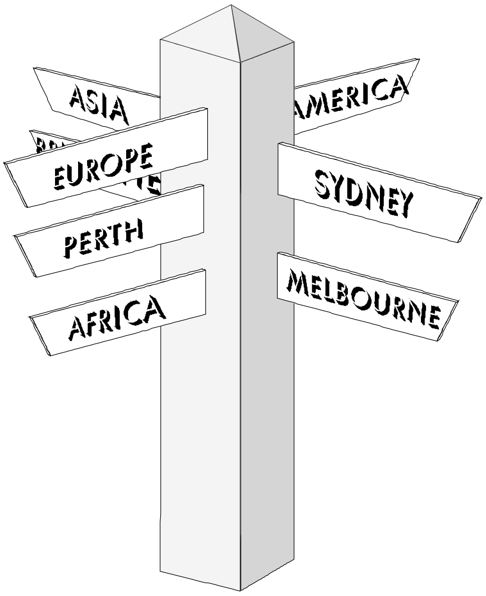

Figure 1-2 The Rationalised Signpost

This benefited the signpost's users, who didn't have to search through lots of different signs to find the one they wanted, and the signpost manager's company, as now, maintenance had been reduced to almost zero.

OK, so this is how road signs work, but is it really relevant to network routing?

Of course it is. Every router has an electronic signpost within it - its forwarding table. Each packet that it forwards, must find a sign telling it which direction to go in, otherwise it will be discarded. A Network Manager is akin to our Signpost Manager and must ensure that there is a suitable sign for every packet that needs to be routed.

By replacing whole groups of signs by a single sign, our Signpost Manager brought together the pointers to many different routes and merged them into a single pointer. In effect, he aggregated those routes - he performed *Route Aggregation*. In fact, he went one stage further. Not only did he bring the routes together, but he also replace the list of individual destinations, be a single common destination name. This procedure is properly known as *Route Information Reduction*.

1.2 Structured Addresses and Routing

From this you may conclude that routers adopt a principle similar to that illustrated in Figure 1-2, and minimise the amount of routing information by collecting routes together and signposting routes to appropriate group addresses. Unfortunate, you would not always be right in making such a conclusion.

For example, in the TCP/IP Internet, the routers implemented by the Internet Service Providers are much more like the signpost in Figure 1-1. There's a sign for every network in the world and, when they run out of space to add new signs, the only answer is to get a bigger signpost. In fact, even this isn't true, because for most Internet Service Providers, there aren't any bigger signposts anymore.

The reason why this is so is twofold. Firstly, the network addresses used in the TCP/IP Internet are rather on the small side at only 32-bits long. Secondly, such addresses have traditionally been allocated to networks without any regard to network topology. The first problem is due to the limited horizons of the early Internet developers. No one at that time thought the Internet would grow so big and a 32-bit address was chosen for engineering reasons (i.e. efficient processing) rather than with future growth in mind. The second problem is simply due to any recognition that there needed to be a way (in network address terms) of forming the structured addresses necessary to move away from the over-crowded signpost.

A Network Address is simply a binary number that uniquely identifies a single host computer on the Internet, However, network addresses are not simply names (like London or Paris) which, on their own tell you nothing about where the addressed location actually is. Network Addresses are first of all names of systems on a network, but they must also be parameters to a routing algorithm that is implemented by every router in an internetwork, and their role as parameters constrains the scope for allocating network addresses.

In our signpost example, the address that we were trying to get to wasn't simply (e.g.) London, but in reality would be a structured address (e.g. 221b Baker Street, London, England, Europe). To find the addressed location, we would consult our first signpost:

- if the signpost is in London, then we start looking for a sign first to Baker Street;
- Otherwise, if the signpost is in England, we look for London;
- Otherwise, if the signpost is in Europe, we look for England;
- and finally, if the signpost is not even in Europe, we look for a sign to Europe.

This is the algorithm we employ to use signposts to help us find our destination. We employ it at every signpost we encounter on our journey and, if they are giving us the right information, we will eventually get to our destination.

In the TCP/IP Internet, a Network Address is similarly structured, but into only two parts. The first part is a unique network identifier and the second part uniquely identifies a Host Computer on the network identified by the first part.

Furthermore, the network identifiers were assigned on a "first come first served" basis. In the electronic signposts that exist in every Internet Router, there has to be a "sign" for every assigned network identifier, pointing along the route to that network. If network identifiers had been assigned (e.g.) that 1 to 100 were in North America, 101 to 200 were in Europe, and so on, then there would be opportunity for the "signposts" within each such router to be rationalised as in Figure 1-2. Within organisations, this is often done, with the Host Identifier split up into an internal (within the organisation) network identifier and a smaller Host Identifier However, at the level of the Internet Service Provider, there is a need to keep track of a route to each assigned network identifier, and this is a serious limitation on Internet growth.

If our electronic signposts are to be rationalised, then Network Addresses must be structured in a way that is much greater than simply Host on Network and so that we can address our systems as (e.g.) *Host* on *internal network*, in *organisation*, attached to *Internet Service Provider*, in *Country* or *Region*. Then, for example, the Routers in an Internet Service Provider (ISP) only need to have "signs" for their users, other ISPs in the same country or region, and an ISP in each other Country or region. The number of such "signs" is then unaffected by the attachment of a new organisation to another ISP i.e. the Internet can grow locally without global impact. This is a necessary condition for an Internet that is scaleable (can always grow bigger). Unfortunately, this is not a realistic proposition with addresses of only 32-bits.

1.3 The Allocation of Structured Addresses

By allocating network addresses arbitrarily (at least on a per network basis), the early developers of the TCP/IP Internet have limited its later growth. Fortunately, for the ATN Internet, these problems were already known by the time that the ATN came to be developed and can thus be largely avoided.

The ATN specifies the use of the Connectionless Network Protocol (CLNP) instead of IP. This has the great advantage of large (variable length) addresses, and the ATN takes advantage of this to specify a 160 bit address format. Although it can be argued that such a long address is less efficient to process than a 32-bit address, 160 bits makes it much easier to ensure that similar network addresses are allocated to networks that are near each other in the ATN Internet, and can therefore be used to improve the overall routing efficiency.

This larger address space allows for a structured allocation of addresses to be made. The address may then be broken up into a number of fields (for the purpose of allocation), which then form a nested hierarchy. For example, in a left to right order, the fields may identify region, country, organisation, site, system. All Systems within a given organisation would than have addresses that share a common prefix and those on the same site also share a common (but longer) prefix. In the ATN, such addresses are known as NSAP Addresses and the prefixes are therefore called NSAP Address Prefixes.

With this approach, similar network addresses, as illustrated in Figure 1-3, imply that the addressed destinations are close together in the topology of the network. Indeed, how far down the address (seen as a bitstring) that the two addresses diverge, can be taken as a metric of closeness.

Indeed, in a scaleable Internetwork, such as the ATN, the Routers operate first by labelling routes with the address prefix(es) common to all destinations along the route, and perform routing simply by comparing destination network addresses against such address prefixes and forwarding each packet along the route labelled with the longest matching address prefix. This is very much like the use of a physical signpost described earlier.

Furthermore, as routing is done by such a simple prefix matching rule, the Routers do not

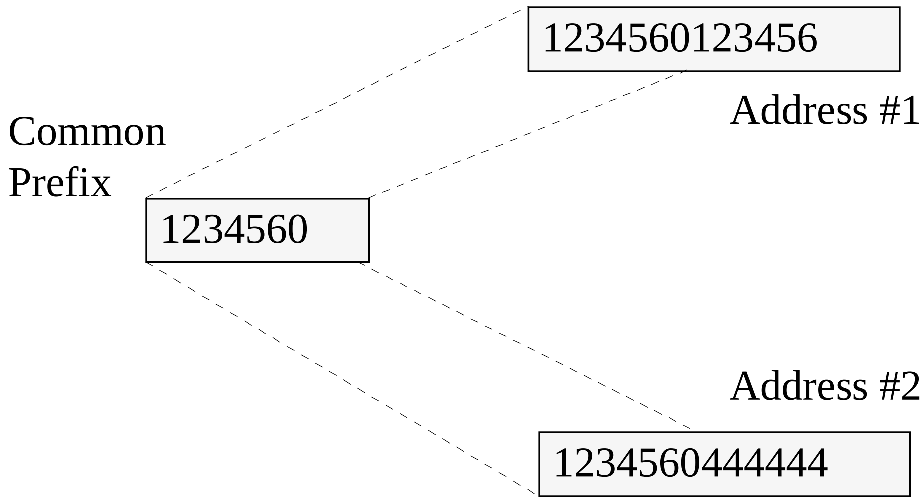

Figure 1-3 Similar Network Addresses

themselves have any real need to know about the structure of the address. The structuring of a network address into a series of fields is therefore only for the purpose of address allocation and not for routing purposes. This is of course different to the way physical signposts are used and represents where our analogy and network routing diverge.

2. Managing Route Information Distribution

In our signpost example, we implicitly assumed that the Signpost Manager had *a priori* knowledge of the routes that were available and applied such knowledge in putting up signs and replacing whole groups of signs by a single sign. Of course, even *a priori* knowledge has to be determine somehow. Our Signpost Manager may have consulted a map prepared by someone else, or he may have journeyed out along each path that led from his signpost, until he found another signpost, and noted down the signs that were on it. Either way, he consulted some information base that was either prepared for him, or which he created himself.

We need to do a similar thing with network routing. Typically a *routing information protocol* is used to enable Routers to either exchange maps, in the form of connectivity information, or to tell each other about the routes they provide. The former approach is generally limited to a local domain (e.g. an organisation), while the latter is generally used for inter-organisation (or more generally known as Inter-Domain) Routing. In a network such as the ATN, the Inter-Domain Routing Protocol (IDRP) is used by Routers to tell each other about the routes they provide.

We also saw our Signpost Manager performing Route Aggregation. in order to first group together related routes, and then Route Information Reduction, in order to replace lists of destination names with a single name. We see the same principles being applied when using IDRP, with the addition of a further procedure - Route Selection.

2.1 Route Selection

In IDRP we don't perform Route Aggregation when creating our electronic signpost. Instead, we do it afterwards, just before we tell neighbouring Routers about the routes we have available. In IDRP, every route that our neighbours tell us about gets put up as a "sign" on the Router's electronic signpost (i.e. gets entered into the Router's Forwarding Table). We can exclude routes, if we don't like the look of them, and will choose between alternatives to the same destination offered by different routers. However, the rule is that each route received, if chosen for use, will get entered into the Forwarding Table.

When we come to tell our neighbours about the routes that we offer - that's when we start aggregating routes, and the first stage in this process is Route Selection.

Route Selection is the process by which the routes that have been selected for inclusion in a Router's Forwarding Table are selected for advertisement to a neighbour Router. The selection process is policy driven - that is rules specified by a Network Manager are applied, in order to determine which routes are to be selected. Furthermore, a different set of rules can be applied for each neighbour router, and therefore, out of the total set of routes present in the Router's Forward Table, a different subset may be advertised to each neighbour router. Indeed, this must be the case, as, at the very least, it doesn't make sense to advertise back to a Router, the routes it advertised to you!

The Route Selection process also selects at least two different categories of routes. There are first the routes that are then to be simply advertised to the neighbour Router, and then there are the routes that are to be first grouped together (i.e. aggregated) before the aggregated route is then advertised to the neighbour Router.

In fact, it is possible to write down a classification for such selection rules. We can define **type 1** rules as being rules that select routes for aggregation i.e. all routes selected by a given type 1 are aggregated before being advertised, and we can define **type 2** rules as being rules that select routes that are individually advertised i.e. each route selected by a type 2 rule is advertised to a neighbour Router.

2.1.1 How IDRP Handles its Routes

As specified in ISO 10747, each router implementing IDRP (properly known as a Boundary Intermediate System or BIS for short), maintains a Routing Information Base (RIB), a Policy Information Base (PIB) and a Forwarding Information Base (FIB). These are illustrated in Figure 2-4.

The IDRP protocol is connection mode and, for each adjacent BIS with which routing information is exchanged, a BIS-BIS connection must exist, and there is required to be a separate RIB data structure to hold all routes received from that BIS. Each such data structure is known as an *adj-RIB-in*. Similarly, for each such adjacent BIS, there is also required to be a separate RIB data structure to hold all routes currently advertised to that BIS. Each such data structure is known as an *adj-RIB-out*.

It should be observed that the contents of a given *adj-RIB-out* ought to be identical to the corresponding *adj-RIB-in* held by the adjacent BIS. To ensure this is the case, the BIS-BIS protocol provides a mechanism to "refresh" an *adj-RIB-in* from the corresponding *adj-RIB-out*. The refresh cycle is periodically performed and ensures that they are identical, thus avoiding any long term persistence of any discrepancies that might have occurred.

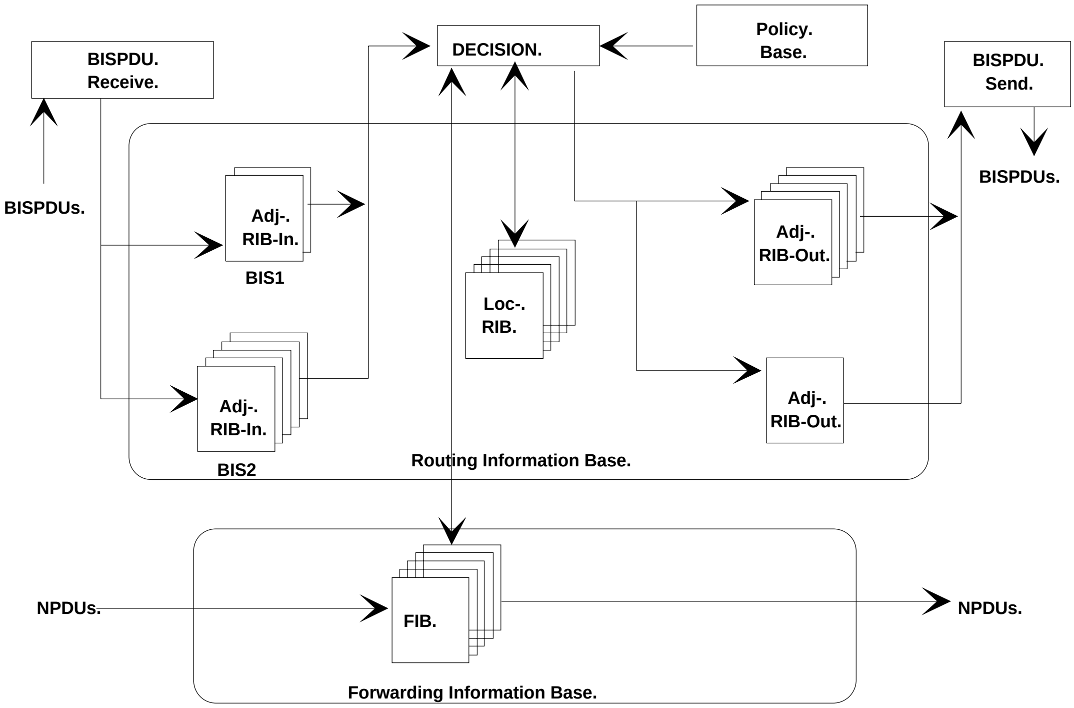

Figure 2-4 IDRP BIS Architecture

The routes received from another BIS are processed by a Routing Decision process, and it is the responsibility of this process to select routes from the adj-RIB-ins for local use and for propagation to other BISs. The so selected routes are recorded in a further data structure, the *loc-RIB*. This holds all currently selected routes, and it is from the loc-RIB that routing information is selected for transfer to the FIB where it is used to support the forwarding of NPDUs, and for the adj-RIB-out for transfer to other BISs.

The Routing Decision Process is formally described as a three phase process, where phases one and two are concerned with the selection of routes and their placement in the loc-RIB, whilst phase three is concerned with the processing of routes in the loc_RIB and maintaining the adj_RIB-out. The PIB may contain rules referenced by each phase of the Routing Decision Process, determining the initial selection of routes, and the aggregation and propagation of so selected routes to other BISs.

It should be noted that IDRP can handled several parallel sets of routes, where each set is distinguished by a set of “Distinguishing Path Attributes” (a RIB_Att). The purpose of this is to permit the distribution of routes according to different Quality of Service metrics and/or security information. In the ATN, two such RIB_Atts are defined: a default RIB_Att under which routes are distributed with no QoS or Security information, and a Security RIB_Att under which routes are distributed with information used to support user driven routing requirements (e.g. to select the use of specific air/ground data link types on a per application basis).

2.1.2 ATN Route Selection

In the ATN, the baseline set of rules for Route Selection and to be implemented by every ATN Inter-Domain Router is specified in section 5.3.7 of the ATN SARPs. In most cases, these rules require the selection of routes for individual advertisement while, in a number of specific cases, rules are required for selecting groups of routes for aggregation into a single route. Typically, a different set of rules is applied for each RIB_Att and for each adjacent BIS.

As has already been discussed, it is possible to classify such rules as Type 1 and Type 2 rules, as illustrated in Figure 2-5.

- All routes selected by a given Type 1 rule are first merged and then the resulting aggregated route entered into the Adj-RIB-out.

- Each route selected by a Type 2 rule is copied as an individual route into the Adj-RIB-Out.

However, analysis of the rules contained in the ATN SARPs, show that there are two types of type 2 rule. This is because there exists a variant on the standard idea of a Type 2 rule, when the application of the rule is conditional on the presence of a specific route in the Adj-RIB-In, that is associated with the same adjacent BIS, as is the Adj-RIB-Out into which the route will be inserted. Furthermore, this route must also be present in the Loc-RIB i.e. if alternatives are available, then this is the preferred one out of the alternatives.

Examples of this variant exist in rules such as:

If the RD is currently advertising the preferred route to all AINSC and ATSC Mobiles, then every route to an AINSC Mobile and an ATSC Mobile that is known to the local RD shall be advertised to this RD, subject only to constraints imposed by any DIST_LIST_INCL and DIST_LIST_EXCL path attributes.

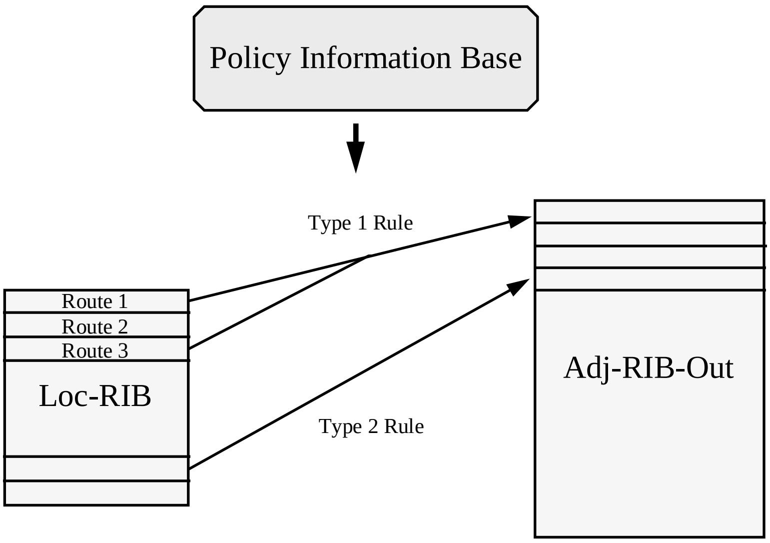

Figure 2-5 Policy Based Route Selection

This type 2 variant is sufficiently different from the definition of Type 2 rules that we have already come across, to make this a clearly identifiable sub-type. The idea of Type 2a and Type 2b rules is therefore introduced, where:

- A Type 2a rule is an unconditional rule for which each route selected by such a rule is copied as an individual route into the Adj-RIB-Out, and

- A Type 2b rule is a conditional rule for which each route selected by such a rule is copied as an individual route into the Adj-RIB-Out, provided that the corresponding Adj-RIB-in also contains a specific route which is also present in the Loc-RIB (i.e. it has been selected for use by the BIS).

The ATN Route Selection rules may therefore be classified as Type 1, Type 2a and Type 2b rules.

If a given route satisfies both Type1 and Type 2 rules then it would also appear correct to assume that the Type 1 rule dominates i.e. that if a route is selected and aggregated with others, then generally it will not be also included in the Adj_RIB_out as an individual route. This appears to be a sensible approach as the whole point of aggregation is to cut down the number of routes being advertised. Special cases can always be explicitly excluded from the aggregation rule.

On the other hand, if a route is selected by more than one Type 1 rule, then while it may be worthwhile issuing a warning to a network manager, it does not appear to be incorrect for it to be included in each such aggregation.

2.1.2.1 Route Selection Criteria

In the ATN SARPs, there are a number of different ways in which routes in the Loc_RIB are selected:

- The NLRI includes/does not include an NSAP Address Prefix in a Network Addressing Domain identified by the same or a shorter NSAP Address Prefix;
- The Route has been originated by or passed through/has not passed through a given Routing Domain;
- Specific Tag Set Types/Tag Set Values are/are not contained in the route’s security path attribute;
- Hop count/Capacity is within defined limits.

Essentially, each rule may be realised as a filter that is being repetitively applied to the routes in the loc_RIB. This filter has both an

- **exclusion condition** (e.g. routes are excluded if they have passed through a Routing Domain known not to accept traffic from the Routing Domain to which the route is being advertised), and an
- **inclusion condition** (e.g. routes are included if they have they required destination NSAP Address Prefix in their NLRI).

2.1.2.2 Types of Filters

Each of the above types of selection rule will therefore consist of a logical expression combining a number of inclusion and/or exclusion filters, selecting or excluding routes on the basis of:

a) RDI of Route originator is a given value

b) Route is via an Routing Domain or RDC with a given RDI

c) Presence of an ATN Security Label tag set and the value of that tag set after the application of a logical ‘AND’ with a mask field equals a specific value.

d) Destination of route either equals or has as its prefix, a given NSAP Address Prefix.

e) Hop Count/Capacity is within defined limits.

The actual syntax of such a logical expression is implementation specific.

2.2 Route Aggregation

Once groups of routes have been selected by Type 1 rules, they will need to be aggregated.

Route Aggregation is formally defined in IDRP as a set of procedures for aggregating two or more routes into a single aggregated routes. This is a purely algorithmic procedure that is fully specified in the standard, and which is applied to every field (path attribute) of the routes in turn.

The Route Aggregation procedures of IDRP also require that the NSAP Address Prefixes that identify the destinations of the routes to be aggregated, are brought together to form a list of NSAP Address Prefixes and which is then the destination of the aggregated route. This is simply the set of NSAP Address Prefixes found in the routes to be aggregated with duplicates removed.

Under the heading "Route Aggregation", It is also possible to replace such lists of NSAP Address Prefixes by a single common prefix, provided one exists. However, this is only possible when there does not exist any NSAP Address Prefix that is not contained in the list and for which the single common prefix is also a prefix. To consider this case, requires the Route Information Reduction procedures.

2.3 Route Information Reduction

Route Information Reduction is the policy based replacement of a list of NSAP Address Prefixes that are contained in a route's destination, by a single NSAP Address Prefix. The list is specified by a Network Manager and is not constrained as to the membership of the list. Such lists are determined taking into account which addresses have actually been allocated, and ensure that routing information can be reduced even when "holes" exist in the addressing plan.

Generally, Route Information Reduction is applied following aggregation. However, it can apply to individually selected routes as well.

There are also two types of Route Information Reduction rule to consider. The first is when a set of NSAP Address Prefixes is replaced by a common prefix only when all members of the set are present in the route. The second type of rule is more relaxed and if any members of a specified set of NSAP Address Prefixes occur in a routes destination, then those that do occur are replaced by the specified NSAP Address Prefix.

The need for these different types of rule is due to the more dynamic nature of network routing compared with our signpost analogy. In a communications network, routes may come and go as systems are switched on and off, or when they fail, and these dynamic changes will need to be reflected in the routing information.

The first type of Route Information Reduction rule always ensures that when part of a network is lost, this is reflected in the routing information exchanged as, such a change will result in one of the routes that makes up the aggregated route being lost, and hence a missing NSAP Address Prefix from the list. In turn, Route Information Reduction is no longer applied and the route is re-advertised correctly advertising what is available.

However, this does result in more routing information being exchanged and more frequently, and there are many occasions when such detail needs to be hidden. For example, if it is known that there is no alternative route to the lost destination, then there is no real harm done in not advertising the change and therefore avoiding the cost in terms of route information distribution. This is the rationale behind the second type of Route Information Reduction Rule.

2.4 Routes with Identical NLRI

Now, it is possible for two routes that have the same NLRI, and which differ only in the security path attribute, to be present in the same Loc_RIB. It is therefore possible for type 2 rules to select more than one route with the same NLRI. It is also possible that as a result of Route Aggregation and Route Information Reduction, two or more routes, to be inserted into the same Adj-RIB-out have the same NLRI.

In all these cases, ISO/IEC 10747 requires that these routes are also aggregated prior to being advertised to an adjacent BIS. This is in order to avoid ambiguity in the advertisement of routing information.

Route Aggregation may therefore have to be applied twice during the phase 3 Routing Decision Process. Once, as a result of the policy based selection of routes to be aggregated (type 1 rules) and, secondly, due to the requirement to aggregate routes with identical NLRI.

Note: the ATN SARPs also specify a simplified version of the Route Aggregation procedures that are called Route Merging. This only applies to a special case when the destinations of the routes to be aggregated are identical. The Route Merging procedures are superseded when proper Route Aggregation procedures are implemented.

2.4.1.1 Update of the Adj-RIB-out

Route Aggregation also adds an additional problem to the general problem of keeping track of which routes are new routes, which have been replaced and which are withdrawn. The IDRP UPDATE BISPDU combination rules permits a route to be withdrawn and replaced by an alternative in a single BISPDU. This is very important in the case of Route Aggregation when a component route is added to or withdrawn from the aggregated route, and the NLRI changes as a result, either directly or because Route Information Reduction is then applied differently In such cases, the revised aggregated routes *replaces* the previous version of the aggregated route. This must happen in a single BISPDU if this is not to affect the other members of the aggregated route i.e. to make these routes unavailable for a period determined by the **minRouteAdvertisementInterval**.

Essentially, the routes in the Adj-RIB-out must be associated with the rule that selected them. Thus when the route that is selected by the rule changes, the proper handling can be determined. Clearly, the handling will depend upon the type of rule.

- A type 1 rule results in one and only one aggregated route. If this changes, then this replaces the previous route;

- A type 2 rule results in one route for each different NSAP Address Prefix found in the NLRI of the selected routes. Such routes may be added to, replaced or withdrawn.

3. Implementation Considerations

3.1 Process Model for IDRP Phase 3

A possible process model for the IDRP Phase 3 Route Decision Process, and which permits Route Information Reduction and Route Aggregation, is illustrated in Figure 3-1. Two PIB data structures are referenced: a list of “Route Selection Rules” and a list of “Reduction Rules”. The former is used for grouping routes together for the purposes of Route Aggregation, while the latter is for determining when Route Information Reduction of NLRI can be performed.

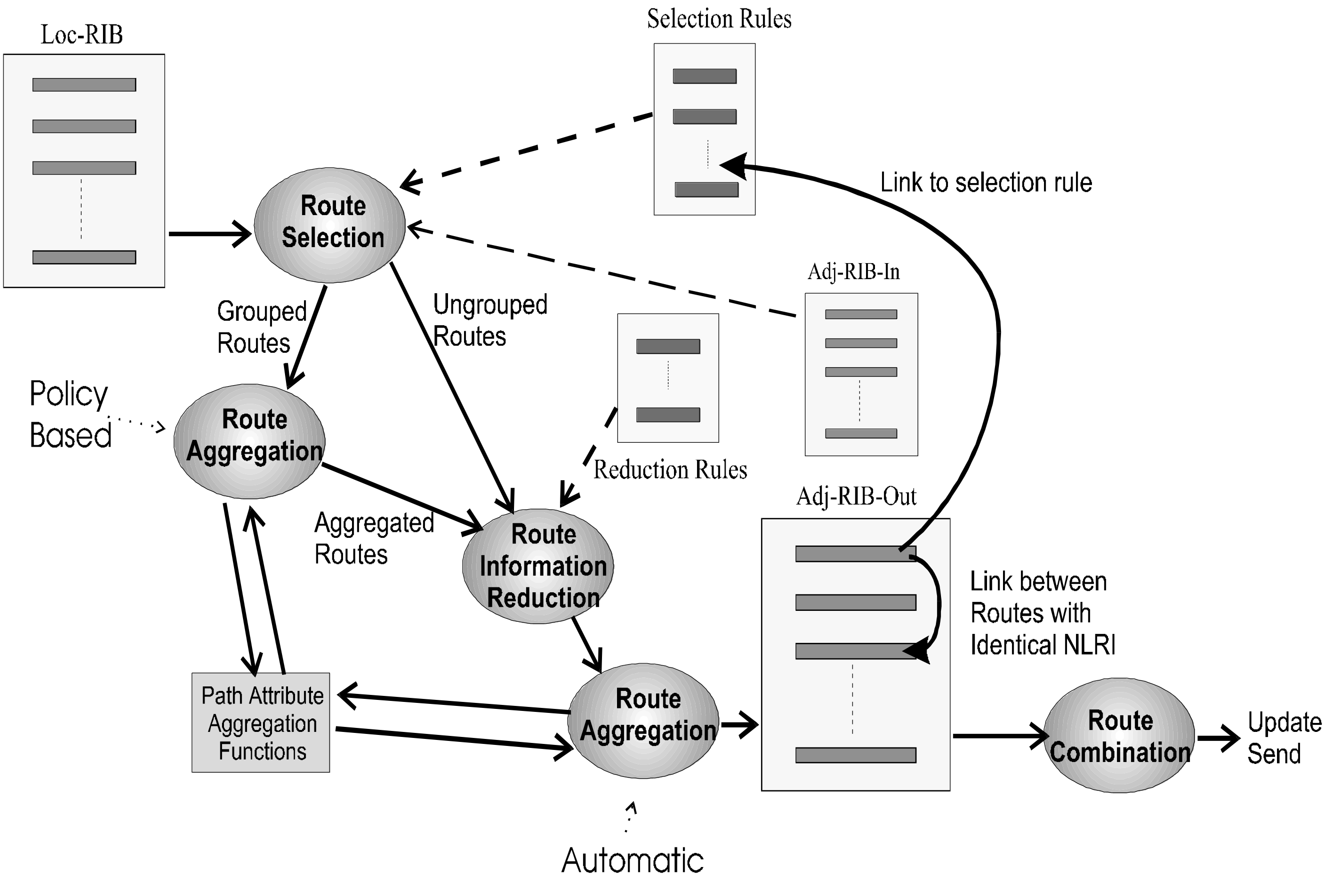

Figure 3-1 Process Model For Route Decision Phase 3

In both cases, it will be necessary for the implementor to define a syntax to enable the text based definition of the rules, so that these data structures may then be created at system start up.

3.1.1 Route Selection

A "Route Selection" process is specified to pass through the Loc_RIB applying first type 1 rules, and then applying type 2a and 2b rules to any routes in the Loc_RIB not selected by a type 1 rule.

3.1.2 Policy Based Route Aggregation

The routes selected by type 1 rules are grouped routes. The routes selected by each type 1 rule form a single group. Each group is then processed by a "Route Aggregation" process to create a single aggregated route for each such group. The aggregation process uses a library of aggregation functions to aggregate each type of path attribute.

Note that some groups of routes cannot be aggregated, even if they have been selected by policy for aggregation. This is because the ISO standard specifically prohibits the aggregation of certain combinations of path attribute. The problem exists for routes that contain:

- DIST_LIST_INCL/EXCL path attributes
- different values of NEXT_HOP
- different values of MULTI_EXIT_DISC.

The outcome, in such cases, is a local matter. However, it is recommended that a deterministic outcome is always ensured.

3.1.3 Route Information Reduction

The remaining routes selected by type 2 rules are ungrouped routes. Both ungrouped routes and the aggregated routes that result from the Route Aggregation process are then passed to a “Route Information Reduction” process. This process inspects the NLRI of each route presented to it and applies the reduction rules to it. The application of a reduction rule will, if the rule is satisfied, result in the replacement of one or more NSAP Address Prefixes in the route’s NLRI, with a single shorter prefix. The rules are applied iteratively until no further reduction can take place.

3.1.4 Automatic Route Aggregation

Once the reduction rules have been applied, the routes are ready to be inserted into the Adj-RIB-out. However, it’s at this point that a check must be made to see if some of these routes have identical NLRI. If they do then they must be aggregated prior to inserting them into the Adj-RIB-out. Note that the same problem may arise, that was discussed above concerning combinations path attributes that cannot be aggregated. In this case, the only solution may be to apply the Route Merging procedures that were specified in the ATN SARPs as a simplified Route Aggregation procedure.

When the routes are inserted into the Adj-RIB-out, they must be linked to the Selection Rule that originally selected it; this is necessary to support the latter processing of the route.

3.1.5 Insertion into the Adj-RIB-out

Prior to inserting the route, the inserting process must check the Adj-RIB-out to see if an existing route is present linked to the same Selection Rule. If this is a type 1 rule, the then new route is marked as replacing the route linked to that Selection Rule. If it is a type 2 or type 3 rule and there is an existing route in the Adj-RIB-out with the same NLRI as the new route, then again the new route is marked as replacing the existing route. Note that in both cases, if the new route is identical to the existing route in both the path attributes it contains and their values then it does not replace the existing route. The existing route may be simply viewed as refreshed.

Indeed, once the phase 3 processes complete, any routes in an Adj-RIB_out that have been neither refreshed nor replaced, must be marked as withdrawn.

3.1.6 Route Combination

Finally, when a route is passed to the Update Send process for advertisement to an adjacent BIS, a “Route Combination” process is required. This will:

a) Ensure that a route withdrawal is always advertised in the same UPDATE BISPDU as the route, if any, that replaces it; and,

b) Ensure that when a route is advertised, it is combined with any routes with the same NLRI, and which are also queued for advertisement to the adjacent BIS.