INTERNATIONAL TELECOMMUNICATION UNION

ITU-T G.7715/Y.1706

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

(06/2002)

SERIES G: TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

Digital terminal equipments – Operations, administration and maintenance features of transmission equipment

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE AND INTERNET PROTOCOL ASPECTS

Internet protocol aspects – Operation, administration and maintenance

Architecture and requirements for routing in the automatically switched optical networks

ITU-T Recommendation G.7715/Y.1706

ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100–G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200–G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300–G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400–G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450–G.499
TESTING EQUIPMENTS	G.500–G.599
TRANSMISSION MEDIA CHARACTERISTICS	G.600–G.699
DIGITAL TERMINAL EQUIPMENTS	G.700–G.799
DIGITAL NETWORKS	G.800–G.899
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900–G.999
QUALITY OF SERVICE AND PERFORMANCE	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DIGITAL TERMINAL EQUIPMENTS	G.7000–G.7999
General	G.7000–G.7099
Coding of analogue signals by pulse code modulation	G.7100–G.7199
Coding of analogue signals by methods other than PCM	G.7200–G.7299
Principal characteristics of primary multiplex equipment	G.7300–G.7399
Principal characteristics of second order multiplex equipment	G.7400–G.7499
Principal characteristics of higher order multiplex equipment	G.7500–G.7599
Principal characteristics of transcoder and digital multiplication equipment	G.7600–G.7699
Operations, administration and maintenance features of transmission equipment	**G.7700–G.7799**
Principal characteristics of multiplexing equipment for the synchronous digital hierarchy	G.7800–G.7899
Other terminal equipment	G.7900–G.7999
DIGITAL NETWORKS	G.8000–G.8999

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation G.7715/Y.1706

Architecture and requirements for routing in the automatically switched optical networks

Summary

This Recommendation specifies the requirements and architecture for the routing functions used for the establishment of switched connections (SC) and soft permanent connections (SPC) within the framework of the Automatically Switched Optical Network (ASON). The main areas covered in this Recommendation include the ASON routing architecture, functional components including path selection, routing attributes, abstract messages and state diagrams.

This Recommendation forms a part of the suite of Recommendations covering the full functionality of the automatic switched transport network (ASTN) and the automatic switched optical network (ASON).

Source

ITU-T Recommendation G.7715/Y.1706 was prepared by ITU-T Study Group 15 (2001-2004) and approved under the WTSA Resolution 1 procedure on 13 June 2002.

Keywords

Automatic Switched Optical Network, Automatic Switched Transport Network, Network-Network Interface (NNI), Network Resources, Path Selection, Routing, Routing Area, Routing Attributes, Routing Control Domain, Routing Controller, Routing Messages, Routing State Machine, User Network Interface (UNI).

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

CONTENTS

Page

ITU-T Recommendation G.7715/Y.1706

Architecture and requirements for routing in the automatically switched optical networks

1 Scope

This Recommendation specifies the requirements and architecture for the routing functions used for the establishment of switched connections (SC) and soft permanent connections (SPC) within the framework of the Automatically Switched Optical Network (ASON). The main areas covered in this Recommendation include the ASON routing architecture, functional components including path selection, routing attributes, abstract messages and state diagrams.

This Recommendation forms a part of the suite of Recommendations covering the full functionality of the automatic switched transport network and the automatic switched optical network. It builds upon the high-level functional requirements and architecture as outlined in ITU-T Recs G.807/Y.1302 (ASTN) and G.8080/Y.1304 (ASON) as the baseline framework for the specification.

This Recommendation is aimed at providing a protocol neutral approach to describe routing for the automatic switched optical networks. Routing messages are transported over a data communication network (DCN). One possible implementation is specified in ITU-T Rec. G.7712/Y.1703.

In order to provide routing service, *a priori* knowledge of the network resources is needed. These resources may be manually provisioned or automatically discovered.

ASON routing has many applications, e.g. traffic engineering, diverse routing, etc. However details of these applications are beyond the scope of this Recommendation.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published.

- ITU-T Recommendation G.805 (2000), *Generic functional architecture of transport networks*.
- ITU-T Recommendation G.807/Y.1302 (2001), *Requirements for Automatic Switched Transport Networks (ASTN)*.
- ITU-T Recommendation G.851.1 (1996), *Management of the transport network – Application of RM-ODP framework*.
- ITU-T Recommendation G.7712/Y.1703 (2001), *Architecture and specification of Data Communication Network (DCN)*.
- ITU-T Recommendation G.8080/Y.1304 (2001), *Architecture of the Automatically Switched Optical Network (ASON)*.
- ITU-T Recommendation M.3016 (1998), *TMN security overview*.
- ITU-T Recommendation X.800 (1991), *Security architecture for Open Systems Interconnection (OSI) for CCITT applications*.

3 Terms and definitions

3.1 This Recommendation uses the following terms defined in ITU-T Rec. G.805:

a) Administrative Domain;

b) Link;

c) Link Connection (LC);

d) Partitioning;

e) Subnetwork;

f) Subnetwork Connection (SNC).

3.2 This Recommendation uses the following term defined in ITU-T Rec. G.851.1:

a) Cluster.

3.3 This Recommendation uses the following term defined in ITU-T Rec. G.7712/Y.1703:

a) Data Communication Network (DCN).

3.4 This Recommendation uses the following terms defined in ITU-T Rec. G.8080/Y.1304:

a) Connection Controller;

b) Federation;

c) Link Resource Manager (LRM);

d) Protocol Controller (PC);

e) Routing Area;

f) Routing Controller (RC);

g) Routing Information Database (RDB);

h) Subnetwork Point Pool (SNPP);

i) Subnetwork Point (SNP);

j) Virtual Private Network (VPN).

3.5 This Recommendation defines the following terms:

3.5.1 node: In the context of this Recommendation, the term node is used to signify a subnetwork or a Routing Area.

3.5.2 Routing Adjacency (RAdj): A logical association between two Routing Controllers.

3.5.3 Routing Control Domain (RCD): An abstract entity that hides the details of the RC distribution. See 5.1 for more information.

3.5.4 Routing Performer (RP): A computational viewpoint object (per ITU-T Rec. G.851.1) that is associated with a routing area and provides an abstraction of the routing service for the routing area.

3.5.5 Shared Risk Group (SRG): A group of resources that share a common risk component whose failure can cause the failure of all the resources in the group.

4 Abbreviations

This Recommendation uses the following abbreviations:

AD Administrative Domain

ASON Automatically Switched Optical Network

ASTN Automatic Switched Transport Network

DCN Data Communication Network

E-NNI External Network-Network Interface

IE Information Element

I-NNI Internal Network-Network Interface

LRM Link Resource Manager

RA Routing Area

RAdj Routing Adjacency

RC Routing Controller

RCD Routing Control Domain

RDB Routing Information Database

RI Routing Information

RP Routing Performer

SNP Subnetwork Point

SNPP Subnetwork Point Pool

SRG Shared Risk Group

UNI User Network Interface

VPN Virtual Private Network

5 ASON routing architecture

The ASON routing architecture supports various routing paradigms listed in ITU-T Rec. G.8080/Y.1304, e.g. hierarchical, step-by-step and source-based. The architecture also abstracts away differences in routing information representation, e.g. link-state, distance-vector, etc. The routing architecture applies after the network has been subdivided into routing areas, and the necessary network resources have been accordingly assigned. The process of subdividing the network into routing areas and assigning network resources is beyond the scope of this Recommendation.

5.1 Fundamental concepts

An operator may choose to subdivide its network based upon specific operator policies, which could include such criteria as geography, administration, technology, etc. The network subdivisions may, by operator decision, be treated as routing areas for the purpose of providing a routing service. Routing areas provide for routing information abstraction, thereby enabling scalable routing information representation. The service offered by a routing area (e.g. path selection) is provided by a Routing Performer (a federation of Routing Controllers), and each Routing Performer is responsible for a single routing area. The RP supports path computation functions consistent with one or more of the routing paradigms listed in ITU-T Rec. G.8080/Y.1304 (source, hierarchical and step-by-step) for the particular routing area that it provides service for. The path computation

functions that may be supported by a RP are based upon the types of information available to it via a Routing Information Database.

Routing areas may be hierarchically contained and a separate Routing Performer is associated with each routing area in the routing hierarchy. It is possible for each level of the hierarchy to employ different Routing Performers that support different routing paradigms. Routing Performers are realized through the instantiation of possibly distributed Routing Controllers. The Routing Controller provides the routing service interface, i.e. the service access point, as defined for the Routing Performer. The Routing Controller is also responsible for coordination and dissemination of routing information. Routing Controller service interfaces provide the routing service across NNI reference points at a given hierarchical level. Different Routing Controller instances may be subject to different policies depending upon the organizations they provide services for. Policy enforcement may be supported via various mechanisms; e.g. by usage of different protocols.

An RC may be implemented as a cluster of distributed entities; such a cluster is called a Routing Control Domain (RCD). An RCD is the abstract entity that hides the details of the distribution internal to the cluster, while providing distribution interfaces with identical characteristics as those of the RC distribution interfaces. The nature of the routing information exchanged between RCDs captures the common semantics of the routing information exchanged between RC distribution interfaces while allowing for different representations within each cluster. The realization of the RCD is implementation specific and outside the scope of this Recommendation.

The relationship between the RA, RP, RC and RCD concepts is illustrated in Figure 1, below.

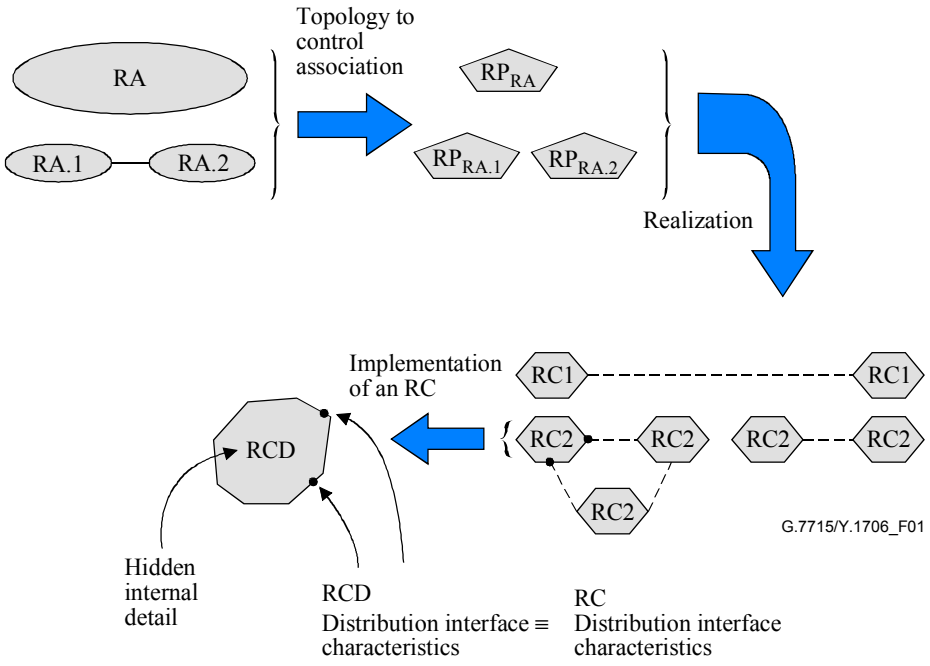

Figure 1/G.7715/Y.1706 – Relationship between RA, RP, RC and RCD

As illustrated above, routing areas contain routing areas that recursively define successive hierarchical routing levels. A separate RP is associated with each routing area. Thus, RP_{RA} is associated with routing area RA, and Routing Performers $RP_{RA.1}$ and $RP_{RA.2}$ are associated with routing areas RA.1 and RA.2, respectively. In turn, the RPs themselves are realized through instantiations of distributed RCs RC1 and RC2, where the RC1s are derived from RP_{RA} and the

RC2s are derived from Routing Performers $RP_{RA.1}$ and $RP_{RA.2,}$ respectively. It may be seen that the characteristics of the RCD distribution interfaces and the RC distribution interfaces are identical[1].

5.2 Routing architecture and functional components

The routing architecture has protocol independent components (LRM, RC), and protocol specific components (Protocol Controller). The Routing Controller handles abstract information needed for routing. The Protocol Controller handles protocol specific messages according to the reference point over which the information is exchanged (e.g. E-NNI, I-NNI), and passes routing primitives to the Routing Controller. An example of routing functional components is illustrated in Figure 2.

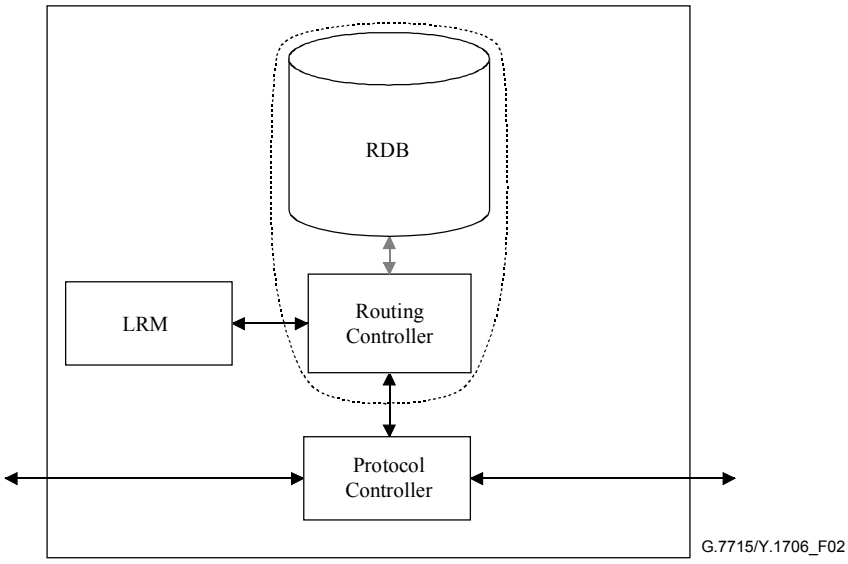

Figure 2/G.7715/Y.1706 – An example of routing functional components

1) Routing Controller – The RC functions include exchanging routing information with peer RCs and replying to a route query (path selection) by operating on the Routing Information Database. The RC is protocol independent.

2) Routing Information Database (RDB) – The RDB is a repository for the local topology, network topology, reachability, and other routing information that is updated as part of the routing information exchange and may additionally contain information that is configured. The RDB may contain routing information for more than one routing area. The Routing Controller has access to a view of the RDB. Figure 2 illustrates this by showing a dotted line around the RC and the RDB. This dotted line signifies the RC (as described in ITU-T Rec. G.8080/Y.1304) as encapsulating a view of the RDB. The RDB is protocol independent.

   NOTE – Since the RDB may contain routing information pertaining to multiple Routing Areas (and hence possibly multiple layer networks), the routing controllers accessing the RDB may share the routing information. This is illustrated in Figure 3 by showing the overlap between the dotted lines.

3) Link Resource Manager – The LRM supplies all the relevant SNPP link information to the Routing Controller. It informs the RC about any state changes of the link resources it controls.

---

[1] Depending upon implementation choice, the number of RCD distribution interface instances need not be the same as that for an RC instance.

4) Protocol Controller – The PC converts the Routing Controller primitives to protocol messages of a particular routing protocol, and is protocol dependent. It also handles protocol specific flow control for the purpose of routing information exchange.

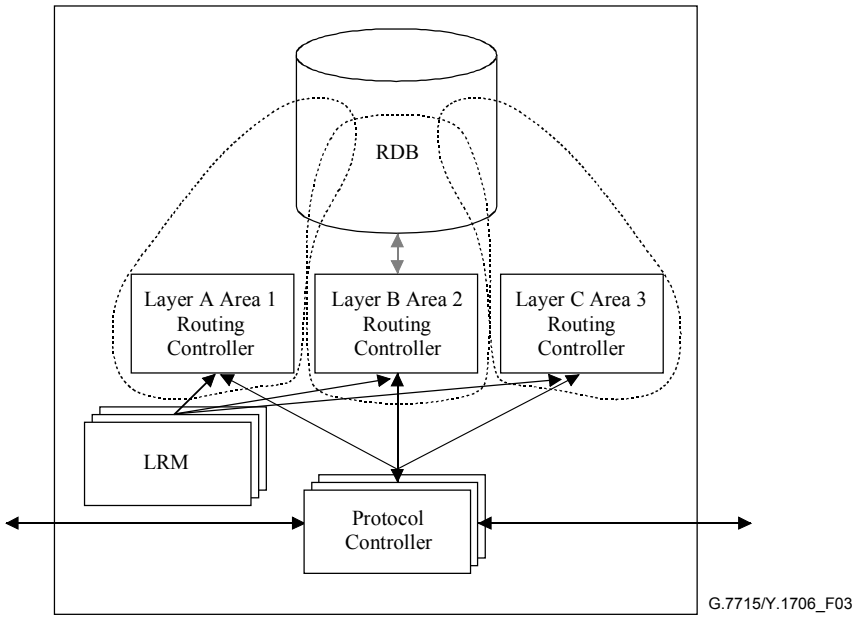

Figure 3/G.7715/Y.1706 – Example relationship between RDB and RCs for several routing areas

5.2.1 Considerations for different protocols

For a given Routing Area, there may be several protocols supported for routing information exchange. The routing architecture allows for support of multiple routing protocols. This is achieved by instantiating different protocol controllers. The architecture does not assume a one-to-one correspondence between Routing Controller instances and Protocol Controller instances. This is illustrated in Figure 4.

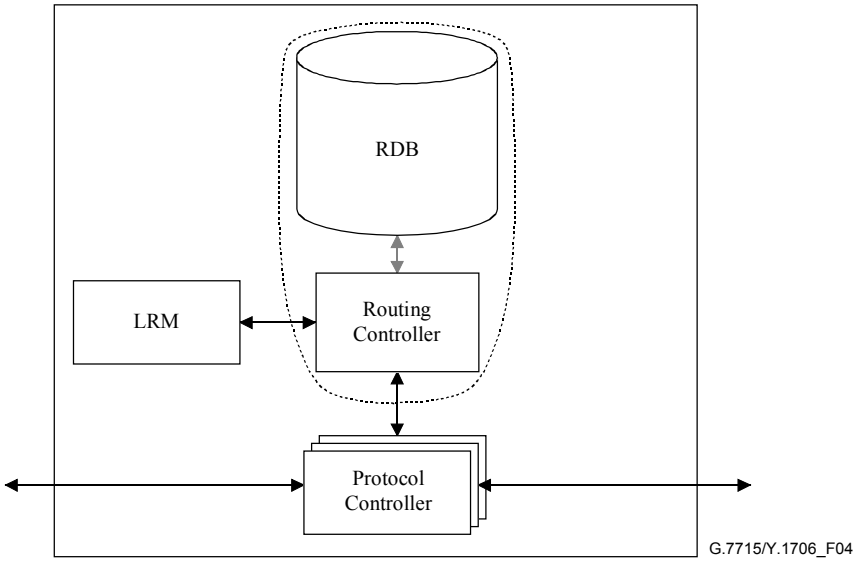

Figure 4/G.7715/Y.1706 – Example architecture for multiple protocols

5.2.2 Considerations for multiple VPNs

A Virtual Private Network is a construct within a single layer network and can be created by:

a) explicitly partitioning network resources for it;

b) sharing common network resources among multiple VPNs.

The routing architecture supports all the above models for VPNs. The explicit partitioning model is supported by defining a view in the RDB for a VPN's RC and populating that view with resources that are not in any other view. This is shown in Figure 5 for the case of VPN3. The shared network resource model is supported by allowing sharing of the RDB by different VPNs, as shown in Figure 5 for the cases of VPN1 and VPN2.

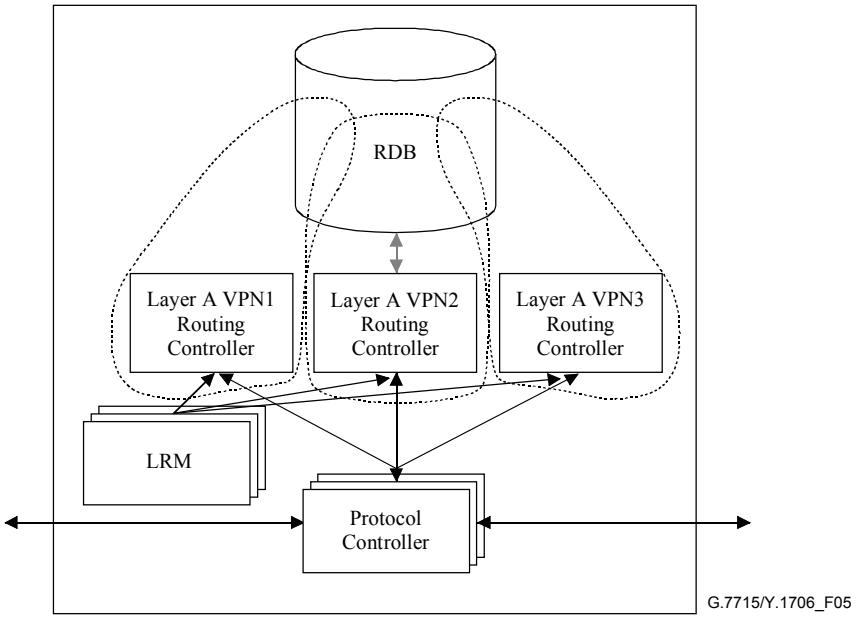

Figure 5/G.7715/Y.1706 – Routing architecture for multiple VPNs

5.2.3 Considerations for policy

Routing policy enforcement is achieved via the policy and configuration ports that are available on the RC component. For a traffic engineering application, suitable configuration policy and path selection policy can be applied to RCs through those ports. This may be used to affect what routing information is revealed to other routing controllers and what routing information is stored in the RDB.

5.3 Routing area hierarchies

An example of a routing area is illustrated in Figure 6 below. The higher level (parent) routing area RA contains lower level (child) routing areas RA.1, RA.2 and RA.3. RA.1 and RA.2 in turn further contain routing areas RA.1.x and RA.2.x.

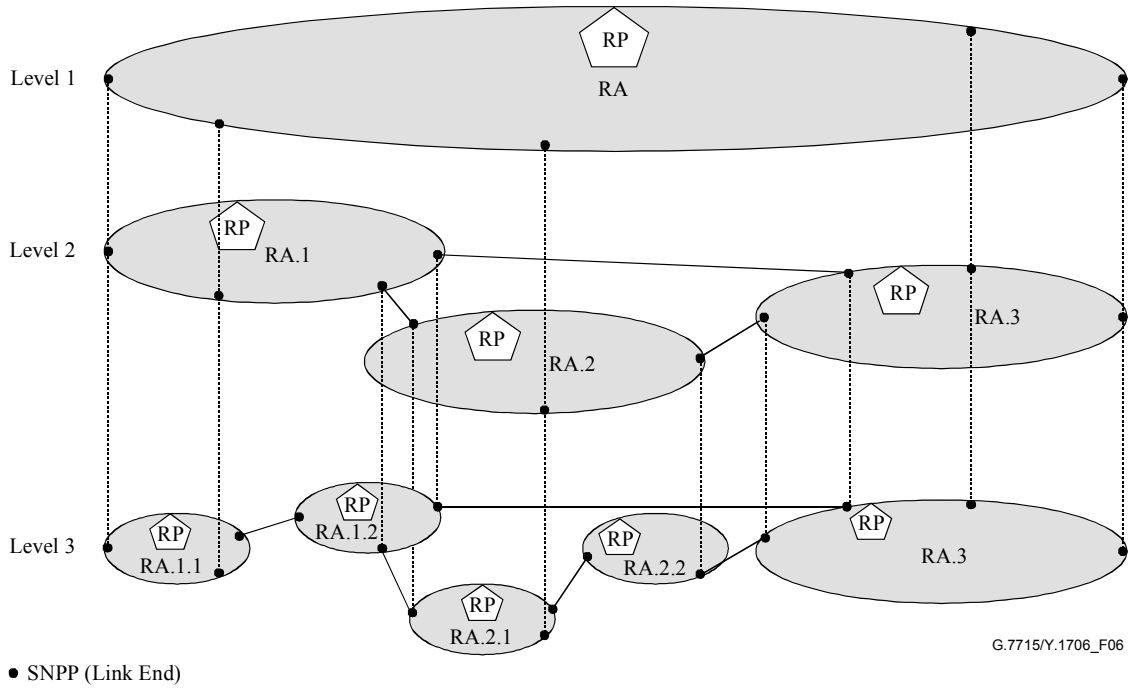

Figure 6/G.7715/Y.1706 – Example of routing area hierarchies

Each routing area has an associated RP that provides the routing service for the routing area at that specific level of the hierarchy. This is illustrated in Figure 7.

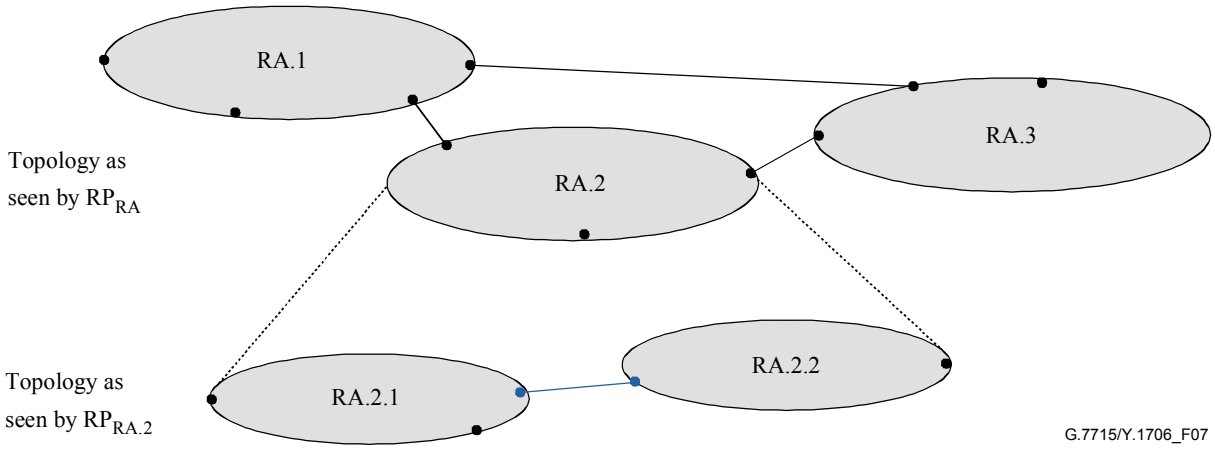

Figure 7/G.7715/Y.1706 – Topology views as seen by RP associated with hierarchical routing areas

5.3.1 Routing performer realization in relation to routing Area hierarchies

The realization of the RP is achieved via RC instances. As described in ITU-T Rec. G.8080/Y.1304, an RC encapsulates the routing information for the routing area, and provides route query services within the area, at that specific level of the hierarchy. In the context of hierarchical routing areas, the realization of the hierarchical RPs is achieved via a stack of RC instances, where each level of the stack corresponds to a level in the hierarchy. Figure 8 depicts the realization of the RPs as a stack of RCs. In the figure, the dashed boxes represent their location within physical elements. For illustrative purposes, the figure shows multiple RC instances collocated in the same physical element; however, this is only an example representation.

At a given hierarchical level, depending upon the distribution choices two cases arise:

– Each of the distributed Routing Controllers could encapsulate a portion of the overall routing information database.

– Each of the distributed Routing Controllers could encapsulate the entire routing information database replicated via a synchronization mechanism.

Note that in either case, the service interfaces of each of the distributed routing controllers are not affected. Distribution also provides the ability to offer multiple service access points within the routing area. There is a need for interaction between the Routing Controllers corresponding to Routing Performers at different levels of the hierarchy (see Appendix I for more information).

NOTE – The special case of a centralized implementation is represented by a single instance of a Routing Controller. (For the purposes of resilience there may be a standby as well.)

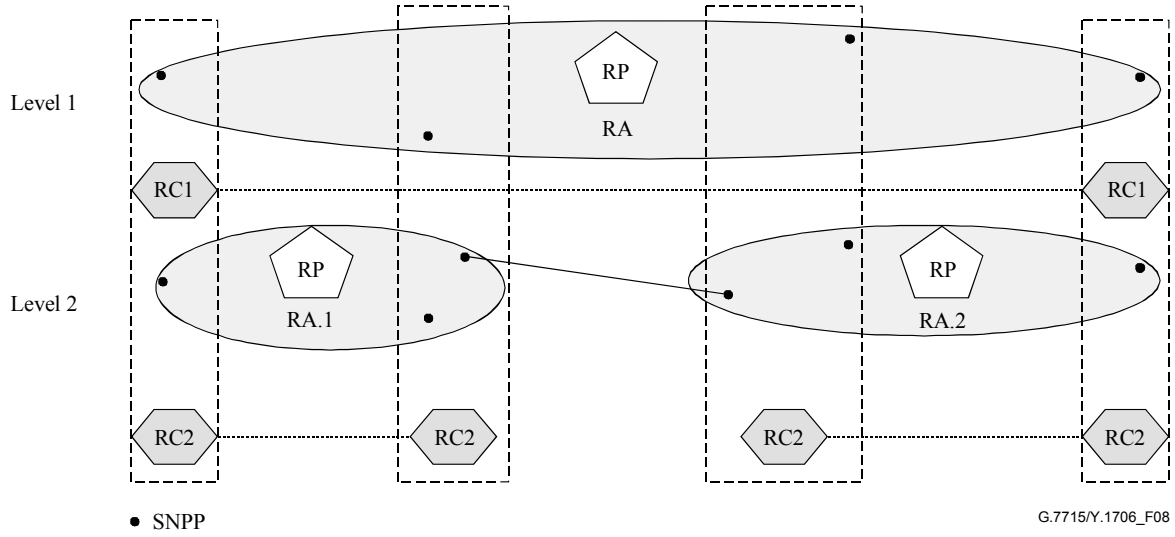

Figure 8/G.7715/Y.1706 – Example Realization of hierarchical Routing Performers

In the context of interactions between Routing Controllers at different levels of the hierarchy, it is important to note that information received from the parent RC shall not be circulated back to the parent RC.

Path selection capabilities have dependencies on the information passed between routing levels and may also require different degrees of subnetwork and link address details. For example, if path selection is not required to return the SNPs of a path, then SNPP level addressing is adequate.

It is possible that a destination will be known to the parent routing area and the child routing area at the same time, but with different paths to the destination. Due to scoping, areas that are closer (lower scope) to the destination will in general have better information to develop a path to the destination than areas that are further away. This is because the child RC always knows all the destinations that are part of its area, or that are part of an area contained within its area. Thus, the child RC is best suited to determine whether it can route directly to the destination.

5.3.2 Correspondence between LRM and RCs for hierarchical routing areas

We use the term "external links" and "internal links" to distinguish between links that are incident upon the routing area and those that are fully encapsulated by a routing area at a given level of hierarchy. Note that external links to a routing area at one level of the hierarchy may be internal links in the parent routing area. LRMs provide the link information to the RCs of the containing routing area. Internal links to a child routing area may be hidden from the parent routing area's view.

Figure 9 shows an example of the association between LRMs (collocated with SNPP) and their corresponding RCs. The LRMs provide the link state information to the Routing Controllers.

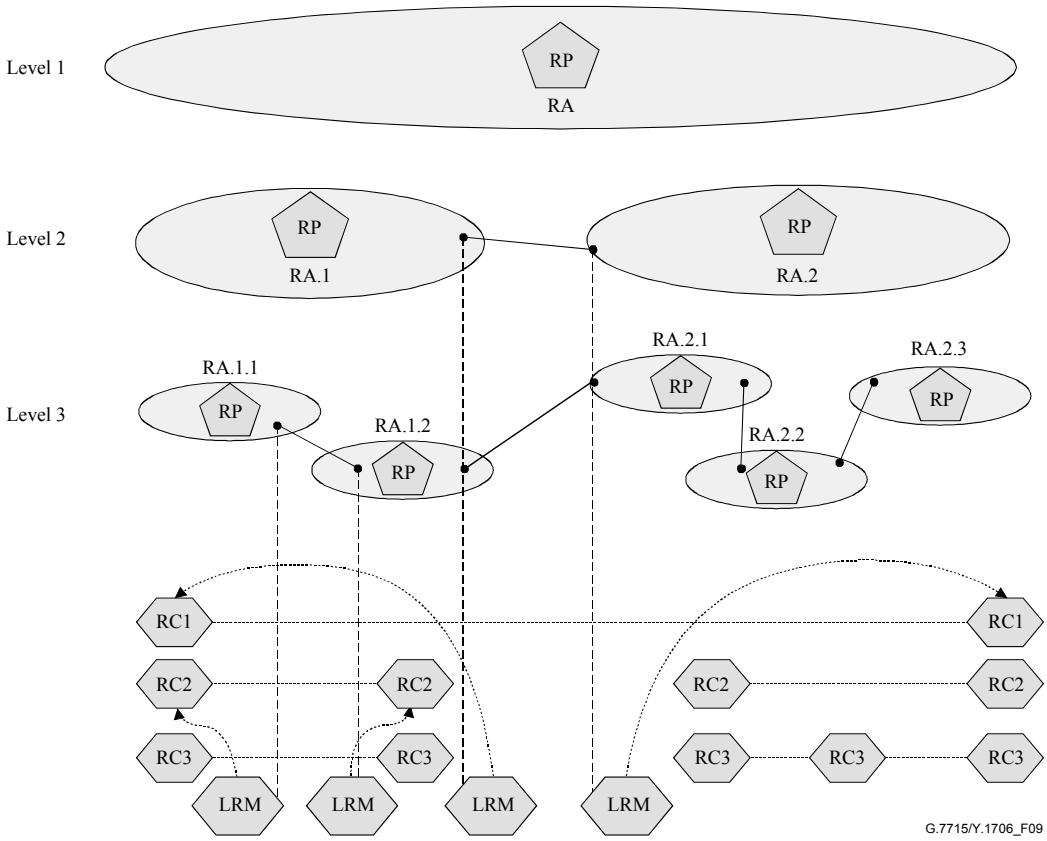

Figure 9/G.7715/Y.1706 – Example association between LRM (corresponding to SNPP) and RC instances

6 ASON routing requirements

ASON routing requirements include architectural, protocol and path computation requirements.

6.1 Architectural requirements

- Information exchanged between routing controllers is subject to policy constraints imposed at the reference points.
- A routing performer operating at any level of hierarchy should not be dependent upon the routing protocol(s) that are being used at the other levels.
- The routing information exchanged between routing control domains is independent of intra-domain protocol choices.
- The routing information exchanged between routing control domains is independent of intra-domain control distribution choices, e.g. centralized, fully-distributed.
- The routing adjacency topology and transport network topology shall not be assumed to be congruent.
- Each routing area shall be uniquely identifiable within a carrier's network.
- The routing information shall support an abstracted view of individual domains. The level of abstraction is subject to operator policy.
- The RP shall provide a means for recovering from system faults (e.g. memory exhaust).

6.2 Protocol requirements

- The routing protocol shall be capable of supporting multiple hierarchical levels.
- The routing protocol shall support hierarchical routing information dissemination including summarized routing information.
- The routing protocol shall include support for multiple links between nodes and shall allow for link and node diversity.
- The routing protocol shall be capable of supporting architectural evolution in terms of the number of levels of hierarchies, aggregation and segmentation of domains.
- The routing protocol shall be scalable with respect to the number of links, nodes and routing area hierarchical levels.
- In response to a routing event (e.g. topology update, reachability update), the contents of the RDB shall converge and a proper damping mechanism for flapping (chattering) shall be provided.
- The routing protocol shall support or may provide add-on features for supporting a set of operator-defined security objectives where required.

  NOTE – Depending on the context of usage of a routing protocol, the overall security objectives defined in ITU-T Rec. M.3016 of confidentiality, data integrity, accountability and availability may take on varying levels of importance. A threat analysis of a proposed routing protocol should address the following items based on ITU-T Rec. X.800; i.e. masquerade, eavesdropping, unauthorized access, loss or corruption of information (includes replay attacks), repudiation, forgery and denial of service.

6.3 Path selection requirements

- Path selection shall result in loop-free paths.
- Path selection shall support at least one of the routing paradigms described in ITU-T Rec. G.8080/Y.1304; i.e. hierarchical, source and step-by-step.
- Path selection shall be able to support a class of routing constraints as described in clause 10.

7 Routing attributes

Information to be disseminated via a routing protocol can be divided between attributes pertaining to links and nodes.

NOTE – The term *node* in the following discussion is used to represent either a routing area or subnetwork depending upon the level of the routing hierarchy at which the path is being computed.

Another important subclass of routing attributes involves those that pertain to resources that can change as a result of connection establishment and deletion procedures. These attributes may require special treatment in networks where connection establishment and deletion dynamics result in fairly frequent changes. For these broad attribute classes, policy and security implications will be considered.

7.1 Node attributes

The main routing attributes to be considered for a node are reachability and diversity-related. A node has complete control over what information it shares. The subclauses below provide a minimal set of information to be shared within a routing area. Depending on the level of trust imposed by policy restrictions at the reference points, the main security implications involve authentication, integrity and non-repudiation.

7.1.1 Reachability attributes

Reachability information provides the set of nodes that are reachable via a given node. This is typically shared via an explicit or summarized list of addresses. The reachability address prefix may include as an attribute the path information from where the reachability information is injected to the destination. Addresses are associated with SNPPs and subnetworks. Operator policy may result in revealing only a subset of reachability information.

7.1.2 Diversity related attributes

The diversity related attributes represent properties of nodes that are used for constrained path selection. One example is the Shared Risk Group (see Appendix II for more information). This attribute, which can be a list of individual node shared risk group identifiers, is used to identify those nodes subject to similar fates.

Another example constraint might be related to exclusion criteria (e.g. non-terrestrial nodes, geographic domains), inclusion criteria (e.g. nodes with dual-backup power supplies).

7.1.3 Other attribute information

Other attributes that could be useful for a node may include additional capability information and the subset of topology information that exists within a child RA.

- Additional capability information is related to the ability to render specific transport services such as protection/restoration capabilities.
- The subset of the topology information that exists within a child Routing Area can be used in diverse transit service offerings, diverse origination and termination services, in traffic engineering, and overall network resource optimization.

Whether any of this information gets shared and the level of detail of this information is subject to policy decisions. Additional attributes (e.g. technology specific attributes) are for further study.

7.2 Link attributes

The minimal set of link attributes to be considered are link state and diversity related. The negotiation of link policy, e.g. glare resolution, is out of scope of the routing function.

7.2.1 Link state

Link State is a triplet comprised of existence, weight and capacity:

- *Existence*

  The most fundamental link attribute is that which indicates the existence of a link between two different nodes in the Routing Information Database. From such information the basic topology (connectivity) is obtained. The existence of the link does not depend upon the link having an available capacity (e.g. the link could have zero capacity because all link connections have failed).

  The choice of disseminating link state information typically goes beyond a policy decision and usually involves a choice between different classes of routing protocols.

- *Link Weight*

  The link weight is an attribute resulting from the evaluation of possibly multiple metrics as modified by link policy or constraint. Its value is used to indicate the relative desirability of a particular link over another during path selection/computation procedures. A higher value of a link weight has traditionally been used to indicate a less desirable path. It may also be used for preventing use of links where the capacity is nearly exhausted by changing the value of the link weights.

- *Capacity*

  For a given layer network, this information is mainly concerned with the number of Link Connections on a link. The amount of information to disseminate concerning capacity is an operator policy decision. For example, for some applications it may suffice to reveal that the link has capacity to accept new connections while not revealing the amount of capacity that is available, while other applications may require the revealing of the available capacity. A consequence of not revealing more information concerning capacity is that it becomes harder to optimize the usage of network resources.

7.2.2 Diversity related attributes

These are similar to the diversity related attribute as described in 7.1.2 except that this relates to links.

7.2.3 Other attribute information

An example of other link attributes is related to availability. Availability is represented in different ways between domains and within domains. Within domains, it could be represented as a specific attribute that indicates the survivability capability of the underlying link. Such information can be particularly valuable if the underlying transport network supports protection and restoration. Between domains this attribute is typically represented in terms of link availability.

8 Routing messages

Functionally, the ASON routing messages may be separated into those pertaining to the maintenance of routing functions such as adjacency maintenance, and those that carry network routing information.

The maintenance messages are exchanged between Protocol Controllers (PC) that have a logical adjacency relationship established between them, either via manual configuration or dynamic establishment. The scope of message exchange is normally confined to the PCs that form the adjacency.

The routing information messages are exchanged between two adjacent Routing Controllers (RCs) and are utilized by the path selection algorithms to calculate and route connection requests across the network. The scope of routing information exchange is normally bounded by its routing area. The messages are propagated via either an incremental, hop-by-hop local exchange or a network-wide flooding mechanism.

Figure 10 shows the routing information messages and events flow between peer RC components. It also indicates the events that are generated upon receipt of protocol messages by the Protocol Controller.

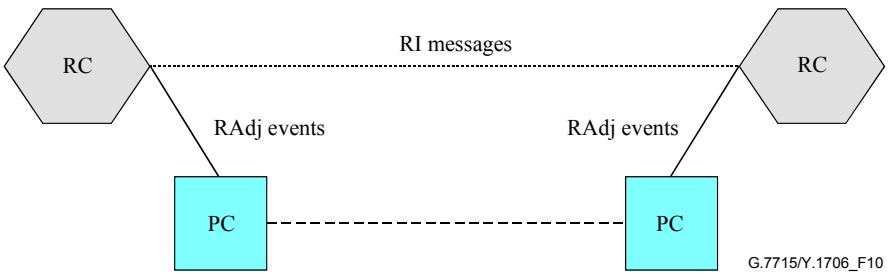

Figure 10/G.7715/Y.1706 – Routing messages and events

8.1 Routing adjacency maintenance

Routing adjacency refers to the logical association between two routing controllers and the state of the adjacency is maintained by the protocol controllers after the adjacency is established. As the adjacency changes its state, appropriate events are sent to the routing controllers by the protocol controllers. The events are used by the routing controller to control the transmission of routing information between the adjacent routing controllers.

The following set of routing adjacency maintenance events are defined:

- RAdj_CREATE: Indicates a new adjacency has been initiated.
- RAdj_DELETE: Indicates an adjacency has been removed.
- RAdj_UP: Indicates a bidirectional adjacency has been established.
- RAdj_DOWN: Indicates bidirectional adjacency has been lost.

8.2 Routing information messages

Routing Information messages are the abstract representation of network routing related information such as node and link attributes as described in clause 7. No routing information is passed over the UNI. Routing information flowing over NNI reference points is subject to the policy constraints at individual NNIs. The routing information exchanged between RCDs over the E-NNI reference point encapsulates the common semantics of the individual RCD information while allowing different representation within each RCD.

Each RC receives and generates routing information messages for the network resources under its direct control and propagates the generated information to adjacent routing controllers. Such a message exchange process will continue until all the RCs' RDBs become stable.

Note that the manner in which information is propagated depends upon the information dissemination mechanism utilized by the routing protocol. The information contained in the routing message varies with the routing protocol used (e.g. link state, distance vector or path vector).

The following generic set of routing messages are applicable independent of the type of the routing protocols.

- RI_RDB_SYNC: These messages help to synchronize the entire routing information database between two adjacent routing controllers. This is done at initialization and may also be done periodically.
- RI_ADD: Once a new network resource has been added, the routing information related to that resource would be advertised using this message in order to be added into the RDB.
- RI_DELETE: Once an existing network resource has been deleted, the routing information related to that resource should be withdrawn from the RDB.
- RI_UPDATE: Once the routing information of an existing network resource is changed, the new routing information related to that resource is re-advertised in order to update the RDB.
- RI_QUERY: When needed, an RC can send a route query message to its routing adjacency neighbour for the routing information related to a particular route.

For the purposes of this Recommendation, this set of messages is specified for the routing information exchange between peer routing controllers.

8.3 Routing exception and error handling

There are many error conditions that may affect the routing process. For example, an RC could receive a corrupted or undefined message during the routing process. In this case, an error message should be generated by the routing protocol to inform the control plane of the error conditions and to enable certain corrective actions. Therefore, a generic notification message is defined as follows to report error or exception condition within the routing domain.

- RE_NOTIFY: This message will be generated when an error or exception condition is encountered during the routing process.

8.4 State diagrams

Three different state machines exist for managing the transmission and reception of Routing Messages.

8.4.1 Information element transmission

The state machine illustrated in Figure 11 and detailed in Table 1 deals with the transmission of routing Information Elements (IE) from a Routing Controller across a routing adjacency to a peer Routing Controller. Throughout the message exchange, it is assumed that the Protocol Controller will provide for the reliable delivery of the transmitted information. One instance of this state machine exists for each Routing Adjacency that is being maintained by the Protocol Controller state machine.

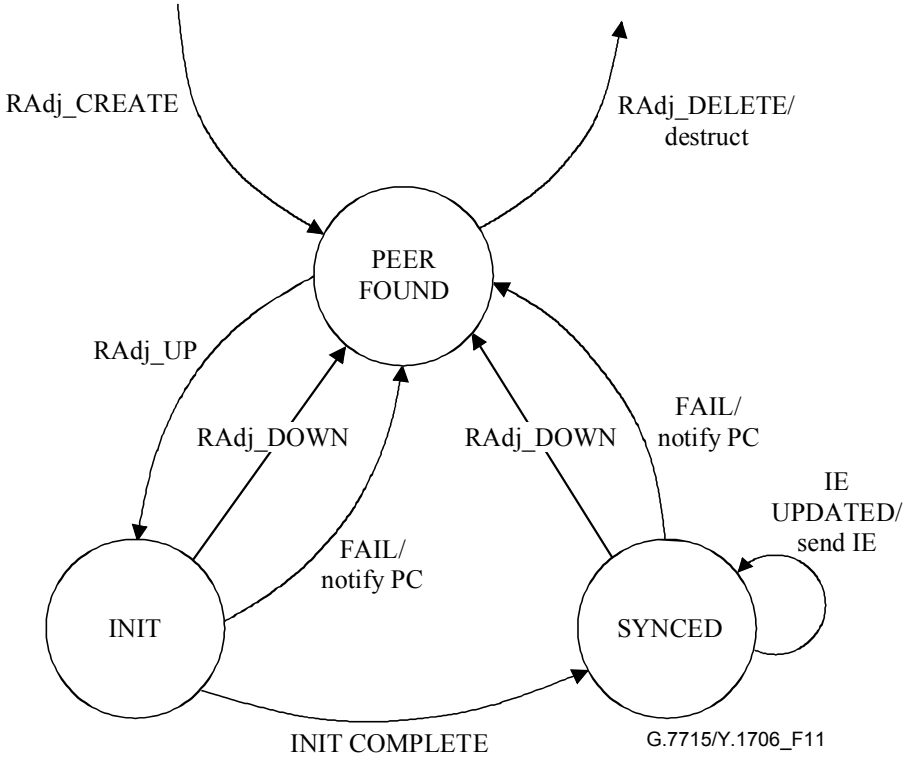

Figure 11/G.7715/Y.1706 – Routing IE transmission state diagram

Table 1/G.7715/Y.1706 – Routing IE transmission state table

State / Event	<Does Not Exist>	PEER FOUND	INIT	SYNCED
RAdj_CREATE	PEER FOUND create state machine	Invalid	Invalid	Invalid
RAdj_UP	Invalid	INIT/start init	Invalid	Invalid
INIT COMPLETE	Invalid	Invalid	SYNCED	Invalid
IE UPDATED	Invalid	Invalid	Invalid	SYNCED/send IE
RAdj_DOWN	Invalid	Invalid	PEER FOUND	PEER FOUND
RAdj_DELETE	Invalid	<Does Not Exist>/ destroy state machine	Invalid	Invalid
FAIL	Invalid	Invalid	PEER FOUND/ notify PC	PEER FOUND/ notify PC

The Routing Controller creates an instance of the state machine when a Protocol Controller identifies a new Routing Adjacency. This is done upon receipt of a RAdj_CREATE event. Initially, the state machine will be in the <PEER FOUND> state. This state exists as a "holding state" until the Protocol Controller identifies the Routing Adjacency as being up. If the Protocol Controller identifies that the routing adjacency no longer exists, then this instance of the state machine is destroyed.

Upon receipt of RAdj_UP event, the state machine will enter the <INIT> state. In this state, the Routing Controller will start the synchronization of the local RDB with the remote RDB.

After the Routing Adjacency has been initialized, the State Machine will enter the <SYNCED> state. While in this state, the local Routing Controller will be notified of changes made to the RDB. When a change occurs, an incremental routing update will be sent to the peer Route Controller.

If the routing adjacency at anytime ceases to be bidirectional, the Protocol Controller sends a RAdj_DOWN event and the state machine will return to the <PEER FOUND> state.

8.4.2 Information element reception

The state machine described in Figure 12 and detailed in Table 2 deals with the reception of Information Elements from a Routing Controller across a routing adjacency to a peer Routing Controller. A single copy of this state machine exists for each Routing Controller.

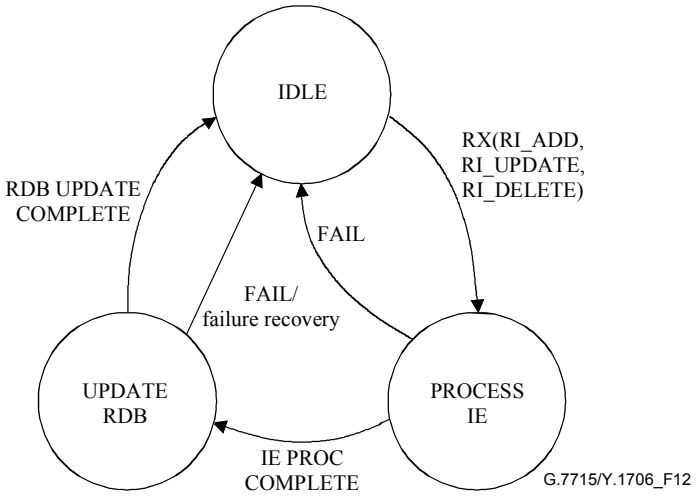

Figure 12/G.7715/Y.1706 – Routing IE reception state diagram

Table 2/G.7715/Y.1706 – Routing IE reception state table

State / Event	IDLE	PROCESS IE	UPDATE RDB
Rx(RI_ADD,RI_UPDATE, RI_DELETE)	PROCESS IE/ start IE processing	Invalid	Invalid
IE PROC COMPLETE	Invalid	UPDATE RDB/ update rdb	Invalid
UPDATE COMPLETE	Invalid	Invalid	IDLE
FAIL	Invalid	IDLE	IDLE/ failure recovery

At the time the routing IE Reception State Machine is initialized, the State Machine will be placed into the IDLE state.

Upon receipt of an RI_ADD, RI_UPDATE, or RI_DELETE message from a peer Routing Controller, the Routing Controller transitions to the <PROCESS IE> state. In this state, the Routing Controller will perform operations on the Information Element to make the information suitable for inclusion into the RDB.

An IE PROC COMPLETE event indicates that the protocol specific processing has been completed, causing the State Machine to submit the IE to the RDB for update based on the Information Element's contents and enters the <UPDATE RDB> state. New information regarding nodes or links will be added to the RDB. Changes to the attributes associated with nodes or links already in the RDB will be handled as an update to the RDB. Likewise, the Information element can direct the Routing Controller to remove a node or link from the RDB.

When the RDB update is complete, an UPDATE COMPLETE event will be received, causing the State Machine to return to the <IDLE> state, where the system will await the reception of another Information Element.

8.4.3 Local information element transmission generation

The state machine illustrated in Figure 13 and detailed in Table 3 deals with Information Elements generated by the RC based on information received from an associated Link Resource Manager. One instance of this state machine exists for each locally generated Information Element.

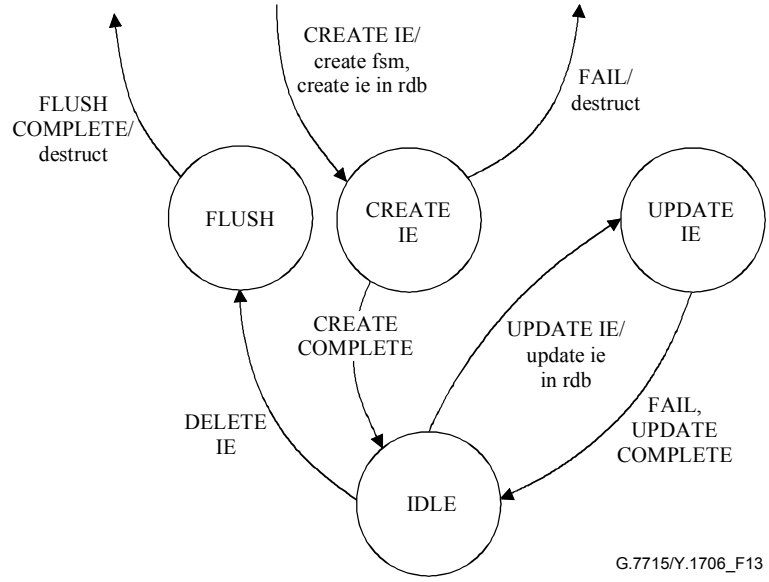

Figure 13/G.7715/Y.1706 – Local information generation state diagram

Table 3/G.7715/Y.1706 – Local information generation state table

State / Event	<Does Not Exist>	CREATE IE	UPDATE IE	IDLE	FLUSH
CREATE IE	CREATE IE/create state machine, create ie in rdb	Invalid	Invalid	Invalid	Invalid
CREATE COMPLETE	Invalid	IDLE	Invalid	Invalid	Invalid
UPDATE COMPLETE	Invalid		IDLE	Invalid	Invalid
UPDATE IE	Invalid	Invalid	Invalid	UPDATE IE/ update ie in rdb	Invalid
DELETE IE	Invalid	Invalid	Invalid	FLUSH	Invalid
FLUSH COMPLETE	Invalid	Invalid	Invalid	Invalid	<Does Not Exist>/ destroy state machine
FAIL	Invalid	<Does Not Exist>/ destroy state machine	IDLE	Invalid	Invalid

As the Routing Controller receives information from an associated Link Resource Manager, the Routing Controller will identify the need to create a new Information Element. As a result, the Routing Controller will create a new instance of the Local Information Generation State Machine, submit the new information element to the RDB, and transition to the <UPDATE IE> state.

When the Information Element has been stored in the RDB, an UPDATE COMPLETE event will be generated. This will cause the State Machine to enter the <IDLE> state, where it will wait for either a request for an update to the Information Element or for a request to delete the Information Element.

When the Routing Controller receives an UPDATE event, the State Machine will send the update information to the RDB, and again transition to the <UPDATE IE> state. As with the creation event, when the RDB has been successfully updated an UPDATE COMPLETE event will be generated, causing the state machine to transition to the IDLE state.

When the Routing Controller receives a DELETE event, the Information Element will need to be deleted from the RDB. Consequently, a flush operation is invoked, and the state machine transitions to the <FLUSH> state.

When the flush is complete, the state machine will receive a FLUSH COMPLETE event, and the Routing Controller will destroy the state machine.

9 Routing message distribution topology

When the Routing Performer for a routing area is realized as a set of distributed Routing Controllers, information regarding the network topology and reachable endpoints needs to be disseminated to, and coordinated with, all other Routing Controllers. The method used to pass routing information between peer Routing Controllers is independent of the location of the source and the user of the information. Consequently, a number of different topologies may be used to pass routing information. The following subclauses provide descriptions of some approaches.

9.1 Congruent topology

Figure 14 shows a routing area containing a network of nodes where all pairs of nodes connected by one or more transport links also have a routing adjacency between the associated Routing Controllers. The resulting topology of routing adjacencies is congruent with the transport network topology. As a result, the nodes connected by the transport link are guaranteed to have visibility to the addresses reachable via the other node. As the transport network approaches a fully connected mesh, the amount of redundant information in circulation increases significantly.

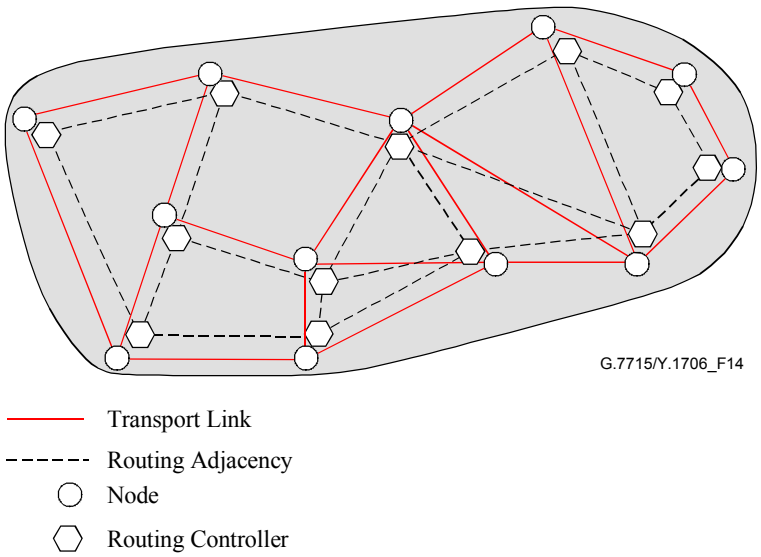

Figure 14/G.7715/Y.1706 – Example of congruent connections

9.2 Hubbed topology using a routing message server

Figure 15 shows routing area containing a network with one or more routing message servers in the network. Each Routing Controller in the network maintains an association with the message server. The message server will in turn pass routing messages received from one Routing Controller to all other Routing Controllers in the Routing Area. Consequently, as Routing Controllers are added to the network, the number of associations to the routing message server scales linearly. Further, since all messages are passed through a common hub point, it is possible for the hub to enforce any policy that exists on the distribution of routing information. However, this approach can result in the transmission of a routing message multiple times on the same DCN link since two different Routing Controllers may be reached across the same DCN link. As mentioned before standby Routing Message Servers can be employed for resilience.

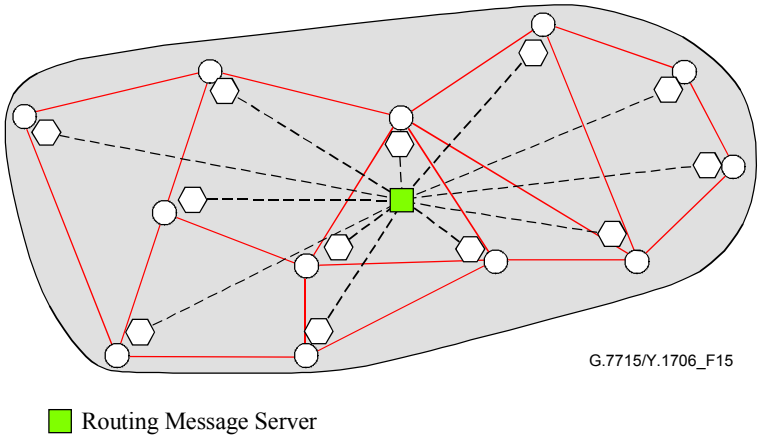

Figure 15/G.7715/Y.1706 – Example of hubbed routing message flows

9.3 Directed topology

Figure 16 shows routing area containing a network where all Routing Controllers in the network forward routing messages according to a pre-provisioned distribution topology. Since the network administrator defines the distribution topology, it can take into account any branching characteristics of the DCN topology, allowing the amount of redundant routing information passed

in the network to be minimised. This distribution topology shall cover all the Routing Controllers in the Routing Area.

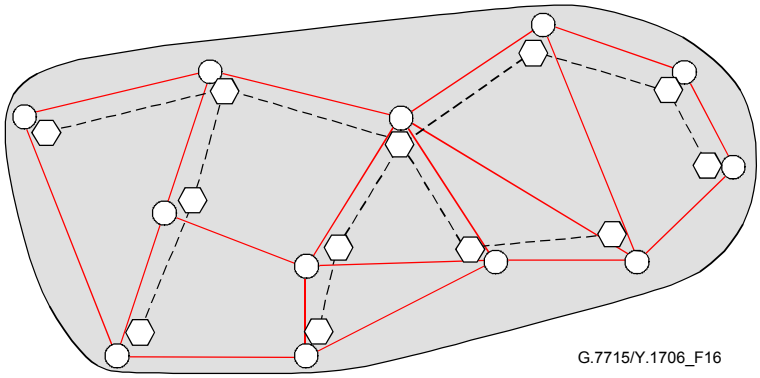

Figure 16/G.7715/Y.1706 – Example of directed topology routing message flows

10 Path selection

ASON path selection is a function that returns a path that a Connection Controller can use as a parameter when signalling a connection. This Recommendation primarily deals with path selection to support connection set-up.

Path selection can be done off-line (route planning by management plane) or on-line in real time (control plane). The choice depends upon the computational complexity, topology information availability, and specific network context. Both off-line and on-line path selections may be provided. For example, operators could use on-line computation to handle a subset of path selection decisions and use off-line computation for complicated traffic engineering and policy related issues such as demand planning, service scheduling, cost modelling and global optimization.

10.1 Inputs to path selection

The inputs to path selection vary depending upon the routing paradigm utilized, examples of which are given in 7.3.2/G.8080/Y.1304. ASON path selection may support a variety of procedures some of which have constraints as input to the path selection request. Examples of constraints are:

- diversity;
- network performance objectives;
- management policies;
- transport layer specific constraints (possibly a link weight metric).

10.1.1 Step-by-step routing

Input to step-by-step path selection typically includes the following factors:

I0 Topology context.

I1 Destination.

I2 Source node.

I3 A set (possibly empty) of constraints that direct what to do when there are multiple possible outputs.

Step-by-step path selection is typically invoked at each node to obtain the next link on a path to a destination. When a source node is supplied as input, it implies that path selection is able to discriminate based on a source/destination pair and not just on destination. That is, a source address is included in the context of making a next hop decision.

Step-by-step path selection is invoked multiple times from different points in the network, and each time the input parameters should be the same. It is also required that the set of path selection instances invoked should collectively result in a path that does not loop back on itself.

I0 represents the place in the network where path selection is invoked and is an important part of the context for step-by-step path selection. For example, step-by-step path selection can be invoked several times with the same I1, I2, and I3 values but return different results if the function is called from different nodes.

10.1.2 Source and hierarchical routing

Source routing and hierarchical routing have similar requirements on path selection. When determining a hierarchical route over subnetworks that are not at the "bottom", the result is the same as a source route with incomplete detail.

Input to source routing path selection typically includes the following factors:

I0 Topology context.

I4 Destination.

I5 Source representing the same, or higher level node, or SNP.

I6 Diversity constraint – Output paths must be diverse from each other.

I7 Inclusion constraints – Links, subnetworks to include in output paths.

I8 Exclusion constraints – Links, subnetworks, or path components to exclude from any output paths.

I9 Minimization metric – This specifies a link metric that the path selection function should minimize for output paths.

The context of source routed path selection is usually a source node for a path. It could also be an intermediate node on a path being calculated. This would typically be at the edge of a routing area where the path selection function for that routing area is invoked to get details on how to cross that routing area.

Hierarchically routed path selection would start at the top of a hierarchy and obtain a sequence of subnetworks through which a path could be found between a given source and destination node. For each of the subnetworks involved, a further path selection is needed that understands the scope of the internal topology of the subnetwork. Conceptually, this recourses until actual links are output by the collective path selection functions. In this case, source routed path selection functions do not have to be identical between two subnetworks.

Inputs I4 and I5 could specify any level of subnetwork, from fabrics to higher level subnetworks.

10.2 Output of path selection

The output of these procedures also varies depending on the routing paradigm. Path selection for source and hierarchical routing are similar in that their outputs are forms of a path (links and nodes), whereas step-by-step routing only requires the next link as its output. In the former case, there are potentially many varieties. This leads to two broad classes of output, characterized in terms of link versus path. Example outputs of this classification include:

O1 Next hop link.

O2 Single path.

O3 Two or more paths

10.3 Routing paradigms and path selection

The routing paradigm determines the required inputs and outputs to the path selection process. Not all combinations of inputs and outputs are useful. Table 4 below shows some combinations that may be of particular interest.

Table 4/G.7715/Y.1706 – Examples of path selection scenarios

Scenario	Inputs	Outputs	Paradigm	Scope of topology required
C1	I0, I1	O1	Step-by-Step. Destination based next hop.	Need view of best output link for a destination address.
C2	I0, I1, I2	O1	Step-by-Step. Source/destination based next hop.	Need view of topology to take full advantage of source context. Could do some discrimination of source with topology from C1 and additional path information.
C3	I0, I4, I5	O2	Source or hierarchical. Source/ destination pair based path. Could be at any level of hierarchy.	Need view of topology at the level of routing hierarchy required.
C4	I0, I4, I5, I8	O2	Source or hierarchical. Output path is routed on a different set of links and nodes with respect to an input path.	Need view of topology at the level of routing hierarchy required.
C5	I0, I4, I5, I7	O2	Source or hierarchical. All links in output path have a common characteristic (e.g. protection)	Need view of topology at the level of routing hierarchy required. Need link characteristic information.
C6	I0, I4, I5, I6	O3	Source or hierarchical. Two output paths are diverse from each other.	Need view of topology at the level of routing hierarchy required.

Appendix I

Information flow between levels of the routing hierarchy

When a connection request is presented to the control plane, it is possible that the local Routing Controller instances do not contain sufficient information to determine how to set up the connection. Address reachability and address discovery is the process by which sufficient information is obtained to determine the first step necessary to set up the connection. Information flow between levels of the routing hierarchy are presented here as illustrative examples. Two arrangements for information flow between routing hierarchical levels are presented in this appendix.

In the first case, Routing Controllers have an interface for interactions between a parent and child RC. During the routing information dissemination time, information is moved up and down the routing hierarchy so that each RC instance is able to adequately resolve any end-point address.

In the second case, resolution is done by recognizing that the local RC does not have sufficient information and forwarding the route query to a parent RC that is presumed to have (or capable of obtaining) the required resolution. This method also has a variant when a Routing Controller has no knowledge of its parent, and resolution is performed by external components.

I.1 Information dissemination related to resolving end-point addresses

Path selection in a routing area at a given layer may be expected to resolve addresses outside of its area. In that case, reachability information from its parent layer does have to be sent downward. This is similar to how many telephone books publish country codes or area codes that are summarized addresses at a higher layer.

When the parent RC provides summarized destination reachability and/or topology to the child, it associates the destination reachability with the child routing area's exit points:

1) A route based on a reachable destination known to the child area because either it is part of the child area, or it is a part of an area contained within the child area shall have the highest preference.
2) A reachable destination that matches a summarized address provided by the parent area to the child area shall have the lowest preference.
3) If the destination of the route request is not known to the child area or to the parent area, then the destination is unreachable.

I.1.1 Parent to child information flow

To support this model, information may be sent from a parent to a child RC to assist path selection in the child RC. The child RC should maintain knowledge of the fact that this information was learned from another level and did not originate at its own level. There are two types of parent information:

1) Reachable addresses – Summarized information about reachable addresses may be sent to lower routing levels to enable their path selection procedures to resolve a destination address to one of their edge SNPPs. Often an association between summarized higher level addresses to addresses owned by the routing area is maintained (e.g. the set of X.75 gateway points through which a particular X.121 Data Network International Code (DNIC) can be reached).
2) Link and Subnetwork – Higher level link and subnetwork information may be sent to lower routing areas. This would enable path selection to make some decisions from the topology of other routing areas. Propagating all higher level information downward has implications on scaling.

I.1.2 Child to parent information flow

Routing information may be also sent from a child RC to a parent RC to be used as local topology input to the parent RC level. Only information originating from the child RC or its children can be sent upward. Routing information received from a parent RC should not be resent back to the parent PC. If insufficient information is sent upward, the higher level path selection procedure may not be able to work.

When a parent RC receives child RC routing information, it may send it downward to other child RCs and/or upward to its own parent RC.

There are two types of child information:

1) Reachable addresses – Addressing information used by routing are routable addresses. This is distinct from public addresses for SNPPs connected to clients, or client names. A mechanism should exist to map those addresses to routable ones, but this is outside the scope of this Recommendation.

Addressing information can be summarized before being sent upward. This is done even if the complete address space is not used by the lower level (i.e. unused addresses or "holes"). More fine grained summaries could be sent from one routing area to its parent level if gaps in an address prefix are to be exposed to the higher level. This is particularly advantageous for step-by-step path selection as a non-existent destination address can be determined sooner.

Address information can also be withheld or excluded. Addresses and/or their summarization cannot be sent to a higher level. This prevents the higher layer from knowing about those addresses. This may be useful for creating routing areas that cannot be used for terminating or transiting connections, but could be used for originating connections.

2) Link and Subnetwork – These entities can be summarized to varying degrees. In complete summarization, a single address representing the whole topology of a routing area can be sent to a higher level. The addresses of SNPPs at the edge of the routing are also sent. In partial summarization a reduced representation of a topology may be sent to a higher level. An example of this is a non-trivial complex node representation in ATMF Private Network-Network Interface (PNNI).

Information detail can also vary according to input requested. For example, if path from a given source to destination is requested, a shortest path tree rooted in that source and containing all destinations would be adequate. If however path from a given source to destination is requested and is to exclude all elements of a given input path, then all links and nodes (subnetworks) in the routing area are needed.

I.2 Information exchange between hierarchical levels for resolving end-point addresses

Path selection in a routing area at a given level may not be expected to resolve addresses outside of its area. In that case, reachability information from its parent level does not have to be sent downward. However, in this case, a route query function provided by the higher layer will be needed so that the lower layer may locate the exit point from the area.

Note that in order to support this method, a limited amount of Child to Parent information transfer is still required.

When hierarchical co-operation is used to identify the exit point for destinations outside of the routing area:

1) The child RC shall first be consulted to develop a path to the destination. If the child RC knows the destination, the path developed by the child RC shall be used. This path shall have the highest preference.

2) When the child RC does not know the destination, the parent RC shall be requested to develop a path to the destination. If the parent RC is able to develop a path, the first link end of the path returned will identify the SNPP used to exit the child routing area. The child RC will next be consulted for a route to the SNPP. The path that is returned by the child RC is then prepended to the path that is returned from the parent RC. This path shall have the lowest preference.

3) If the parent RC is unable to develop a path, then the destination is unreachable.

(NOTE – A parent is allowed to contact its parent also).

Appendix II

Shared Risk Group

This appendix introduces and discusses the concept of Shared Risk Group (SRG) that is useful in support of computing path diversity between a pair of source and destination nodes in the circuit-switched networks.

II.1 Path diversity

In circuit-switched networks, a connection traverses a path that is alternating between nodes and links. One of the most common requirements is for multiple parallel circuits between the same pair of source and destination node to be routed over diverse network resources (links and nodes) across the network. Diversity is a relationship between circuits, and the objective of diverse routing is to eliminate the single point of failure that could potentially fail more than circuits due to a single network resource failure such as a fibre cut or switching node failure.

At the abstract level, there are two types of topological path diversities: link disjoint and node disjoint[2]. Two circuits are said to be link disjoint if they do not share a common link as a potential single point of failure. Similarly, two circuits are said to be node disjoint if they do not share a common node as a potential single point of failure except the originating and the terminating nodes. Note that in the graph theoretic sense node disjoint implies link disjoint; however, in real networks geographic information needs to be taken into consideration for achieving this.

The above diversity definition is based on the topological relationship of the network resource in terms of links and nodes. Such topological diversity can be extended based on the risk-sharing model discussed below.

II.2 Network resources and risk sharing

An end-to-end connection utilizes both link and node resources and network resource failure includes fibre cuts, switching node crashes, central office fires, etc. A set of network resources tend to fail together when they share the same fate due to geographic proximity, which can be used to define shared risk groups. Some common examples include:

- Fibre Sharing: All the wavelengths carried in the same pair of fibres in a WDM-based Optical Transport System.
- Cable/Conduit Sharing: All the fibres within a cable or contained in the same conduit due to physical construction constraints.
- Right of Way Sharing: All the fibre routes that share the same right of way. A Right of Way is land where the network operator has the right to install conduit or fibre cable.

Networks are often built sharing common facilities or rights of way with other industries or utilities. Similarly, in a metropolitan network, different carriers might lease duct space in the same conduit or may lease fibre facilities from the same carrier. In all the above cases, the fault hazard on the common infrastructure is the shared risk.

---

[2] One may also have link and node disjoint.

The risk-sharing concept should not be limited to the constraints imposed by the physical network structure. Administrative policy-based risk sharing groups are also useful. Some examples below illustrate the concepts:

- Two circuits that might be considered diverse for one application might not be considered diverse for another. Diversity is usually thought of as a reaction to interoffice route failures.
- High reliability applications may require taking into account other types of failures. Office outages do occur, although less frequent than route failures, fires, power outages and floods. Many applications require node diversity. In other cases, it may only require power diversity from the same office.
- Shared Rings: Many applications allow "diverse" circuits share a SONET ring-protected link; presumably they would allow the same for optical layer rings.
- Disaster Area: Earthquakes and floods can cause failures over an extended area. For diversity purposes, all the network resources located within an earthquake or flood zone can be viewed as sharing the same risk.
- Some networks may tolerate higher risks. For example, some network operators might consider two fibre cables in a heavy duty concrete conduit as having a low chance of simultaneous failure, hence link diverse. They might view two fibre cables buried on opposite sides of a railroad track as being diverse because there is a minimal chance of single backhoe disrupting both of them.

II.3 Shared Risk Group (SRG)

We can extend the common node and link diversity to the general Shared Risk Group, which can affect nodes, links or both. Specifically, we refer to SRG-diversity as opposed to node/link-diversity; the latter being a special case of the former.

We define a *Shared Risk Group (SRG)* as a group of elements that share a common risk, whose failure can cause the failure of all the elements in the group.

For example, all the fibre links that go through a common conduit in the ground belong to the same SRG group because the conduit is a shared risk component whose failure, such as a cut, will cause all the fibres embedded in the conduit to break.

A SRG identifier is often defined for each shared risk group for identification purpose. The SRG concept can be used to define link-disjoint path diversity. Two data paths are SRG disjoint if no two links or nodes on the two paths belong to the same SRG under consideration.

The SRG concept works well for a flat network topology and can be extended to hierarchical network. A flat network can be partitioned into domains that consist of a set of nodes and associated links. The partition can be based on topological, technological or administrative reasons. The links between domains are the links between the nodes in different domains and each domain can be further partitioned into sub-domains. Such partition can be repeated until the granularity of the domains reach a certain level, thereby generating a network hierarchy. In the hierarchical network, a domain is a logical node that has a set of ports corresponding to the incoming and outgoing links.

It is important to note that the SRG concept not only defines a shared risk group of nodes and links, but also defines the preference level in terms of path selection by considering the risk level and path quality. SRG is essential in supporting hierarchical routing in which a domain level path can be calculated first and further expanded within each domain.

II.4 SRG implications for routing

Dealing with diversity is an unavoidable requirement for routing in the circuit-switched transport network. It requires dealing with the SRG constraints in the routing process, but most importantly requires additional state information for the SRG relationships.

At present, most SRG information cannot be self-discovered. Indeed, in a large network it is very difficult to maintain accurate SRG information. It is particularly challenging whenever multiple administrative domains are involved; for example, after acquisition of one network by another, because there normally is a likelihood that there are diversity violations between domains. It is unlikely that diversity relationships for SRGs based on measures other than geographic relationships (e.g. latitude, longitude coordinates for links and nodes) will be used at any time in the near future.

There is considerable variation in what is meant by acceptable diversity, so an SRG could be characterised by 2 parameters:

- Type of Compromise: Examples would be shared fibre cable, shared conduit, shared Right of Way, shared optical ring, shared office without power sharing, etc.
- Extent of Compromise: For compromised outside plant, this would be the length of the sharing.

A Constrained Shortest Path First (CSPF) algorithm could then penalize a diversity compromise by an amount dependent on these two parameters.

Note that globally consistent SRG information is not always available across multiple control domains, even within a single carrier.

The mapping between links/nodes and different SRGs is in general defined by network operators based on SRG policies and rules. Since SRG information is not yet ready to be discovered by a network element and does not change dynamically, it might not be advertised with other resource availability information by network elements. It could be configured in some central database and be distributed to or retrieved by the nodes, or advertised by network elements at the topology discovery stage.

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE AND INTERNET PROTOCOL ASPECTS

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	**Y.1700–Y.1799**
Charging	Y.1800–Y.1899

For further details, please refer to the list of ITU-T Recommendations.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	**Transmission systems and media, digital systems and networks**
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communications
Series Y	**Global information infrastructure and Internet protocol aspects**
Series Z	Languages and general software aspects for telecommunication systems

Geneva, 2002