



INTERNATIONAL TELECOMMUNICATION UNION

**ITU-T**

**G.7715.1/Y.1706.1**

TELECOMMUNICATION  
STANDARDIZATION SECTOR  
OF ITU

(02/2004)

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, INTERNET PROTOCOL ASPECTS  
AND NEXT GENERATION NETWORKS

Internet protocol aspects – Operation, administration and  
maintenance

---

**ASON routing architecture and requirements for  
link state protocols**

ITU-T Recommendation G.7715.1/Y.1706.1

---

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 TESTING EQUIPMENTS	G.450–G.499
TRANSMISSION MEDIA CHARACTERISTICS	G.500–G.599
DIGITAL TERMINAL EQUIPMENTS	G.600–G.699
DIGITAL NETWORKS	G.700–G.799
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.800–G.899
QUALITY OF SERVICE AND PERFORMANCE – GENERIC AND USER-RELATED ASPECTS	G.900–G.999
TRANSMISSION MEDIA CHARACTERISTICS	G.1000–G.1999
DIGITAL TERMINAL EQUIPMENTS	G.6000–G.6999
General	G.7000–G.7999
Coding of analogue signals by pulse code modulation	G.7000–G.7099
Coding of analogue signals by methods other than PCM	G.7100–G.7199
Principal characteristics of primary multiplex equipment	G.7200–G.7299
Principal characteristics of second order multiplex equipment	G.7300–G.7399
Principal characteristics of higher order multiplex equipment	G.7400–G.7499
Principal characteristics of transcoder and digital multiplication equipment	G.7500–G.7599
Principal characteristics of transcoder and digital multiplication equipment	G.7600–G.7699
<b>Operations, administration and maintenance features of transmission equipment</b>	<b>G.7700–G.7799</b>
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.7715.1/Y.1706.1**

## **ASON routing architecture and requirements for link state protocols**

### **Summary**

This Recommendation provides architecture and requirements for a link-state realization of ITU-T Rec. G.7715/Y.1706. It identifies protocol-neutral requirements for hierarchical link state routing derived from ITU-T Recs G.8080/Y.1304 and G.7715/Y.1706 in a distributed environment. The distribution of the architectural components is defined in ITU-T Rec. G.8080/Y.1304, and this Recommendation is one realization of the ASON routing architecture.

### **Source**

ITU-T Recommendation G.7715.1/Y.1706.1 was approved on 22 February 2004 by ITU-T Study Group 15 (2001-2004) under the ITU-T Recommendation A.8 procedure.

## 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.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

## 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.

© ITU 2004

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

## CONTENTS

	<b>Page</b>
1 Introduction .....	1
2 References.....	1
3 Terms and definitions .....	2
4 Abbreviations.....	2
5 Architecture framework for ASON link state routing.....	2
6 Routing hierarchy .....	3
7 Identifiers for routing.....	5
7.1 Identifier spaces.....	5
7.2 Routing component identifiers .....	5
7.3 Name space interaction.....	6
7.4 Name spaces and routing hierarchy.....	6
7.5 SNPP name components.....	6
8 Routing within a hierarchy .....	6
8.1 Routing information exchange .....	6
8.2 LRM to RC communications.....	9
8.3 Configuring the hierarchy and information flow.....	9
8.4 Configuring RC adjacencies.....	9
9 Routing attributes .....	9
9.1 Principles .....	9
9.2 Taxonomy of attributes.....	10
9.3 Common advertisement information.....	10
9.4 Node attributes.....	10
9.5 Link attributes.....	11
Appendix I – Use of SNPP aliases in multilevel routing hierarchies .....	13
Appendix II – Multiple RC communications over PCs .....	14
Appendix III – Example scenarios for RA segmentation, aggregation and evolution of the network hierarchy .....	15
III.1 Segregation – Insertion of a node in an RA that is "full" .....	16
III.2 Merging two RAs .....	16
III.3 Renaming an RA .....	17
III.4 Insertion – Different routing protocol added to a portion of a network .....	18



# ITU-T Recommendation G.7715.1/Y.1706.1

## ASON routing architecture and requirements for link state protocols

### 1 Introduction

ITU-T Recs G.807/Y.1302 and G.8080/Y.1304 specify the requirements and architecture for a dynamic, automatically switched optical network (ASON) in which connection services are provided using a control plane. ITU-T Rec. G.7715/Y.1706 specifies the requirements and architecture for the ASON routing functions used for the establishment of switched connections and soft permanent connections.

This Recommendation provides protocol-neutral requirements for a hierarchical link state routing protocol derived from ITU-T Recs G.8080/Y.1304 and G.7715/Y.1706 in a distributed environment. The distribution of the architectural components is defined in ITU-T Rec. G.8080/Y.1304, and this Recommendation is one realization of the ASON routing architecture. The support of multilevel hierarchical routing is a key function within this instantiation of ITU-T Rec. G.7715/Y.1706.

The requirements for link-state routing protocols supporting control plane communication (i.e., SCN) is not within 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. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- ITU-T Recommendation G.803 (2000), *Architecture of transport networks based on the synchronous digital hierarchy (SHD)*.
- ITU-T Recommendation G.805 (2000), *Generic functional architecture of transport networks*.
- ITU-T Recommendation G.807/Y.1302 (2001), *Requirements for the automatic switched transport networks (ASTN)*.
- ITU-T Recommendation G.872 (2001), *Architecture of optical transport networks*.
- ITU-T Recommendation G.7712/Y.1703 (2003), *Architecture and specification of data communication network*.
- ITU-T Recommendation G.7713/Y.1704 (2001), *Distributed call and connection management (DCM)*.
- ITU-T Recommendation G.7715/Y.1706 (2002), *Architecture and requirements for routing in the automatically switched optical networks*.
- ITU-T Recommendation G.8080/Y.1304 (2001) and Amendment 1 (2003), *Architecture of the automatically switched optical network (ASON)*.

### 3 Terms and definitions

<b>Terminology</b>	<b>Definition</b>
Connection Point	See ITU-T Rec. G.805
Link Resource Manager	See ITU-T Rec. G.8080/Y.1304
Protocol Controller	See ITU-T Rec. G.8080/Y.1304
Routing Area	See ITU-T Rec. G.8080/Y.1304
Routing Controller	See ITU-T Rec. G.8080/Y.1304
Routing Database	See ITU-T Rec. G.7715/Y.1706
Routing Performer	See ITU-T Rec. G.7715/Y.1706
Shared Risk Groups	See ITU-T Rec. G.7715/Y.1706
Signalling Communication Network	See ITU-T Rec. G.7712/Y.1703
Subnetwork Point	See ITU-T Rec. G.8080/Y.1304
Subnetwork Point Pool	See ITU-T Rec. G.8080/Y.1304
Termination Connection Point	See ITU-T Rec. G.805

### 4 Abbreviations

This Recommendation uses the following abbreviations:

CC	Connection Controller
CP	Connection Point
IE	Information Element
LRM	Link Resource Manager
PC	Protocol Controller
RA	Routing Area
RC	Routing Controller
RDB	Routing Information Database
RP	Routing Performer
SCN	Signalling Communication Network
SNP	Subnetwork Point
SNPP	Subnetwork Point Pool
SRG	Shared Risk Groups
TCP	Termination Connection Point

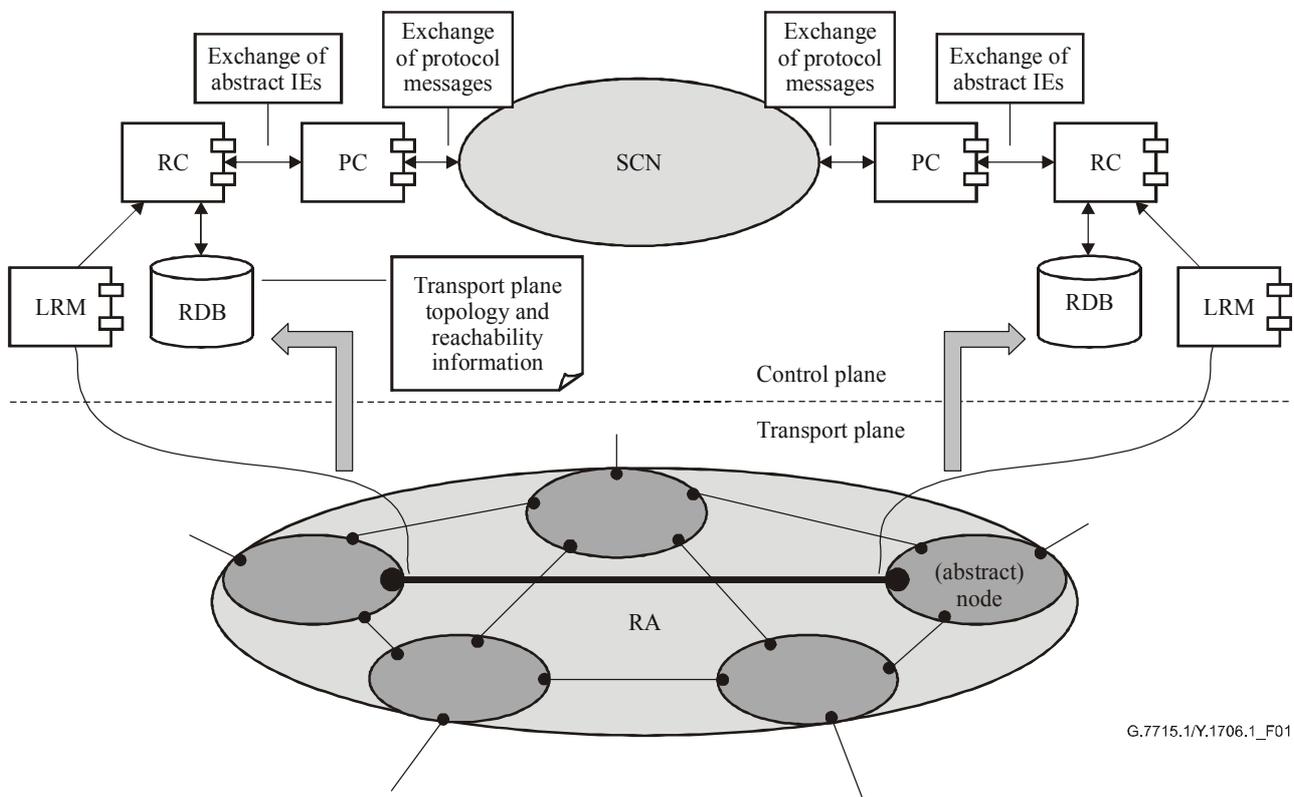
### 5 Architecture framework for ASON link state routing

The architectural framework of ITU-T Rec. G.7715/Y.1706 provides a starting point for further architectural analysis. Figure 1 illustrates a Routing Area (RA) in which the Routing Performer (RP) is instantiated via multiple Routing Controllers (RCs) that belong to the same RA. In this example, two RCs are shown for simplicity. The interaction of functional and G.8080/Y.1304 architectural components related to routing for this RA, described in ITU-T Rec. G.7715/Y.1706, are illustrated. No distinction is made between nodes that may have further internal details

(i.e., abstract nodes) and those that cannot be decomposed any further.

As described in ITU-T Rec. G.8080/Y.1304, the RC components receive link state information from their associated Link Resource Managers (LRMs) regarding SNPP links, and store this information in the Routing Information Database (RDB). The RDB is replicated at each RC within the same RA, and may contain information about multiple network layers in the transport plane. Whenever the state of an SNP changes, the LRM informs the corresponding RC which, in turn, updates its associated RDB. In order to assure RDB synchronization, the RCs cooperate and exchange routing information. In this context, the physical realization of these RC communications is via a particular link state routing protocol. This is represented by the protocol controller (PC) component in the G.8080/Y.1304 architecture, and the protocol messages are conveyed over the SCN as described in ITU-T Rec. G.7712/Y.1703. The PC may convey information for one or more network layers (see Appendix II).

Path computation functions may exist in each RC, on selected RCs within the same RA, or may be centralized for the RA. Path computation on one RC is not dependent on the RDBs in other RCs in the RA. If path computation is centralized, any of the RDBs in the RA (or any instance) may be used. Path computation algorithms are not for standardization and are outside the scope of this Recommendation.



G.7715.1/Y.1706.1\_F01

**Figure 1/G.7715.1/Y.1706.1 – Example of functional and architectural component interactions within an RA in a single level**

## 6 Routing hierarchy

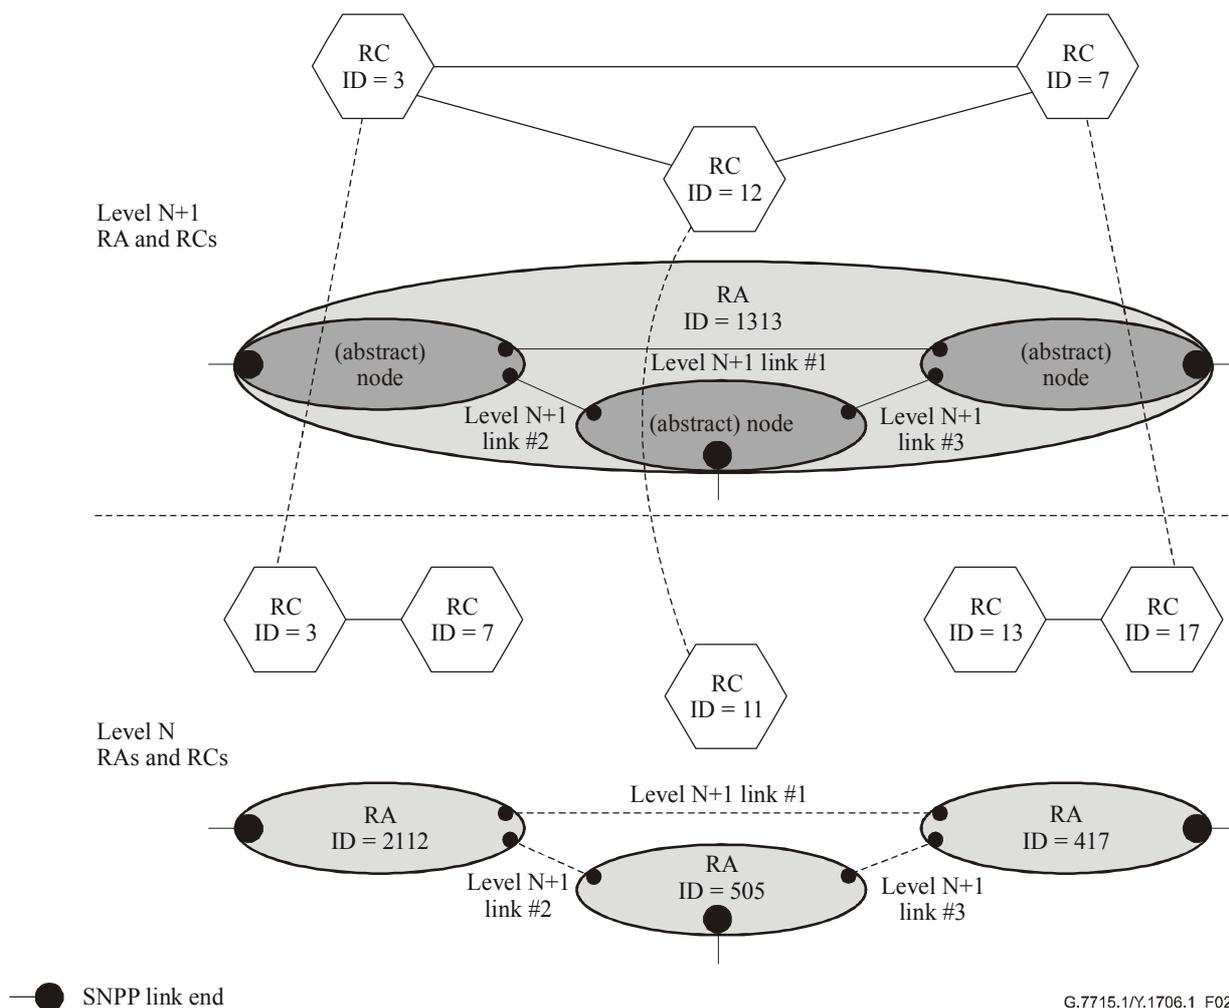
Routing areas allow for information abstraction, enabling information hiding and scalable routing information representation. Except for the single RA case, RAs are hierarchically contained and a separate routing performer is associated with each routing area in the routing hierarchy. It is possible for each level of the hierarchy to employ different routing performers that support different routing paradigms. Routing areas contain routing areas that recursively define successive

hierarchical routing levels. The number of hierarchical levels to be supported is implementation-specific and outside the scope of this Recommendation.

In a multilevel routing hierarchy, it is necessary to distinguish among RCs within a level, and RCs at different levels of the routing hierarchy. Before any pair of RCs establishes communication, they must verify they belong to the same RA. An RA identifier (RA ID) is required to provide the scope within which the RCs may communicate. In order to distinguish between RCs within the same RA, an RC identifier (RC ID) is required; the RC ID must be unique within its containing RA.

NOTE – RA IDs may be associated with a transport plane name space whereas RC IDs are associated with a control plane name space.

An illustrative example of a multilevel routing hierarchy where RC identifiers within one RA are reused within another RA, is provided in Figure 2. In this particular topology, the RC IDs at hierarchy level N+1 overlap with those used within some of the different level N RAs.



**Figure 2/G.7715.1/Y.1706.1 – Example network where RC identifiers within one RA are reused within another RA**

In the process of managing an ASON network, it is anticipated that the containment relationships of RAs may need to change, motivated by unforeseen events such as mergers, acquisitions, and divestitures. Consistent with ITU-T Rec. G.7715/Y.1706, the routing protocol shall be capable of supporting architectural evolution in terms of number of hierarchical levels, as well as aggregation and segmentation of RAs (see illustrative examples in Appendix III). To facilitate this, RA IDs need

to be provisioned to be unique within an administrative domain. The link-state routing protocol is not expected to automatically initiate and/or execute these operations.

## 7 Identifiers for routing

All ASON components represent abstract entities, as described in 7.1/G.8080/Y.1304. The ASON routing component has identifiers whose values are drawn from several identifier spaces. Issues that affect routing protocol requirements include maintaining separation of identifier spaces, understanding what other components use the same spaces that routing uses, and understanding what mappings are needed between spaces.

### 7.1 Identifier spaces

There are three categories of identifiers used for ASON routing: transport plane names, control plane identifiers for components, and SCN addresses.

- Transport plane names describe G.805 resources and are defined in ITU-T Rec. G.8080/Y.1304.
  - SNPP names give a routing context to SNPs and are used by the control plane to identify transport plane resources. However, they do not identify control plane components but represent a (G.805) recursive subnetwork context for SNPs. Multiple SNPP name spaces may exist for the same resources. Note must be taken that the topology of the control network shall not be assumed to be congruent with the transport network topology (ITU-T Recs G.807/Y.1302 and G.7715/Y.1706).
  - UNI transport resource names are used to identify transport resources at a UNI reference point if they exist. They represent clients in (G.8080/Y.1304) access group containers and may be disseminated by RCs.
- Control plane identifiers for G.8080/Y.1304 components may be instantiated differently from each other for a given ASON network. For example, one can have centralized routing with distributed signalling. Separate identifiers are thus used for:
  - Routing Controllers (RCs);
  - Connection Controllers (CCs).

Additionally, components have Protocol Controllers (PCs) that are used for protocol-specific communication. These also have identifiers that are separate from the (abstract) components like RCs.

- SCN addresses enable control plane components to communicate with each other via the SCN as described in ITU-T Rec. G.7712/Y.1703.

### 7.2 Routing component identifiers

Identifiers are required for ASON link-state routing components. These include the following:

- RC and PC identifiers: These are derived from the control plane identifier space. The RC ID is used to identify the entity that creates the routing announcement. The PC ID is used to identify the entity that is involved in the protocol exchange.

NOTE – An implementation may choose to use the same identifier value for both RC and PC identifiers. If it is assumed that the same identifier value is always used, a protocol may be optimized by using a single protocol information element for both the PC and RC identifiers.
- PC (associated with the RC) address for communication over the SCN: These are derived from the SCN address space.
- Transport resource identifiers for resources represented by RCs: These are derived from the SNPP name space. The identifier format used within a routing function is implementation-specific and outside the scope of this Recommendation.

It should be noted that, in order to maintain functional separation among the different ASON routing components, identifier spaces should be independent from each other, i.e., it should be possible to change the SCN addresses used for communication between the PCs without affecting the routing adjacency between peering PCs. This separation, however, does not mean that identical formats cannot be used. For example, an IPv4 address format may be used by multiple name spaces.

### **7.3 Name space interaction**

The SNPP name space is used by routing and signalling functions. Common SNPP name space and formats should be used for interactions between routing and signalling functions. For example, the path computation function of an RC must return a path to a connection controller (CC) in a common SNPP name space and representation. Otherwise, the CC may not understand the path.

### **7.4 Name spaces and routing hierarchy**

ITU-T Rec. G.8080/Y.1304 does not restrict the number of SNPs corresponding to a CP. This means that there can be multiple SNPP name spaces for the same subnetwork. An important design consideration is whether one or multiple SNPP name spaces are used. The following options exist:

- Use a separate SNPP name space per level in a routing hierarchy. This requires a mapping to be maintained between each level, e.g., a translation function performed via a directory service. Mapping functions (e.g., directory services) are outside the scope of this Recommendation. A scenario for this is described in Appendix I.
- Use a common SNPP name space for all levels in a routing hierarchy. As described in ITU-T Rec. G.807/Y.1302, a hierarchical naming format may be used, supporting aggregating addresses in a hierarchical manner. This provides one approach towards supporting hierarchical routing.

### **7.5 SNPP name components**

An SNPP name consists of either one or a set of RA names, an optional subnetwork name, and link contexts. SNPP names are used by RC interfaces to identify transport plane resources. An RC would include SNPP names in information sent to other RCs, as well as information stored in its associated RDB. The result of path computation is also returned in the form of SNPP names.

In an SNPP name, the RA name space represents the scope of an RC. The RA ID format can be derived from any address-space global in scope, including IPv4, IPv6, and NSAP addresses.

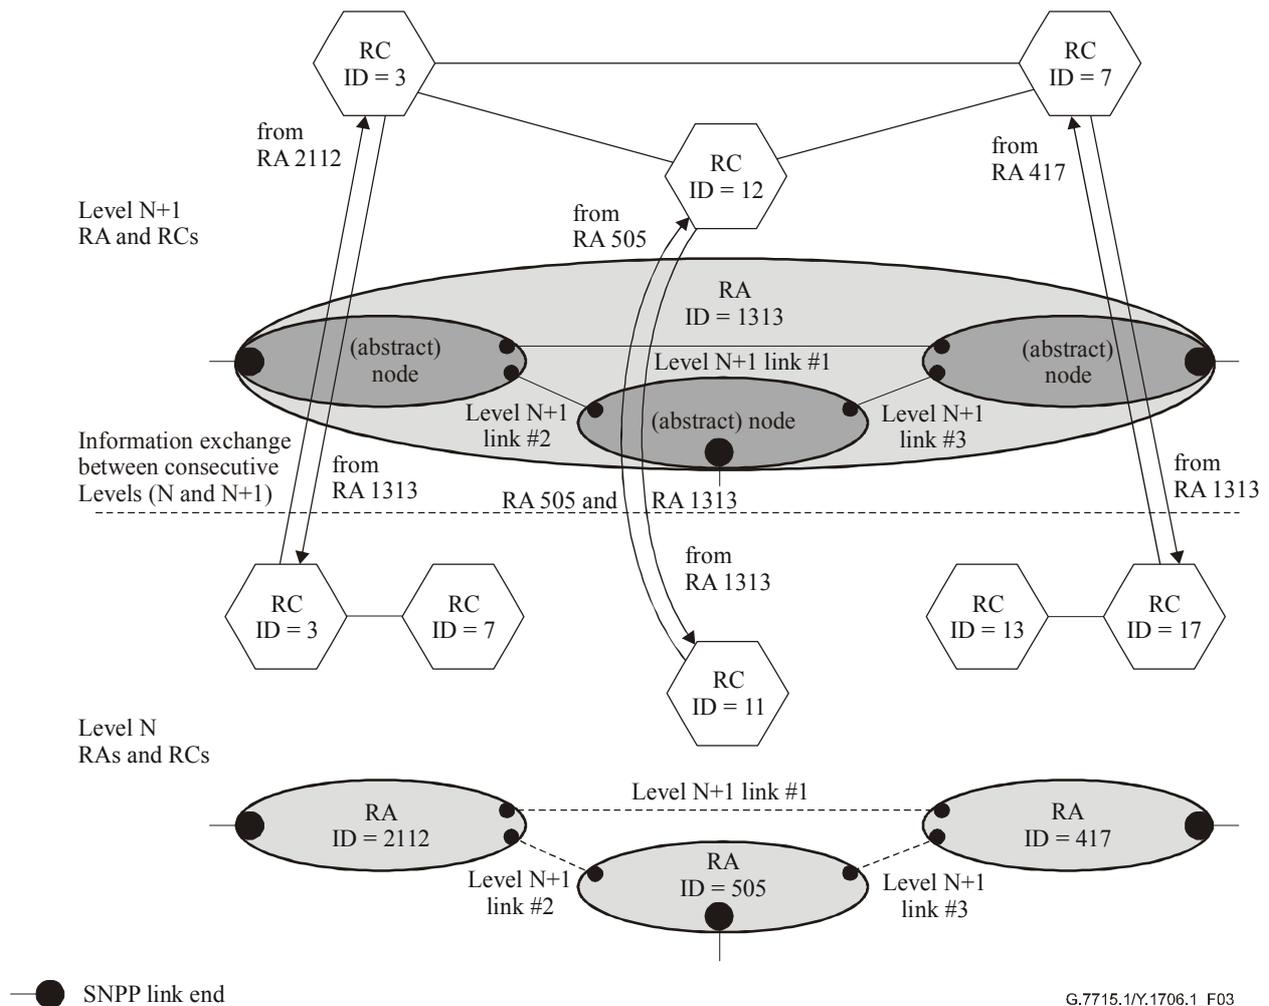
As described in ITU-T Rec. G.7715/Y.1706, a node may represent both an abstraction of an RA or a subnetwork. The node ID is the subnetwork ID that exists in an SNPP name. All node IDs advertised within an RA shall be allocated from a common name space for that RA. It should be noted that node IDs and RC IDs generally do not have a one-to-one relationship and may be allocated from different namespaces.

## **8 Routing within a hierarchy**

In this clause the exchange of routing information between hierarchical routing levels is detailed.

### **8.1 Routing information exchange**

Routing information may be exchanged across different levels of the routing hierarchy. Figure 3 illustrates an example of how such information exchange may occur between two adjacent levels, i.e., level N+1 and N, where level N represents the RAs contained by level N+1. The links connecting RAs may be viewed as inter-RA links, and the links representing connectivity within an RA may be viewed as intra-RA links.



**Figure 3/G.7715.1/Y.1706.1 – Example of information exchange between level N and level N+1 RCs**

In Figure 3, it should be noted that the physical location of, e.g., RC 11 and RC 12, their relationship and their communication protocol is outside the scope of this Recommendation.

NOTE – No assumption is made regarding how RC 11 and RC 12 communicate.

Information exchange between an RC, its parent, and its child RCs, may include reachability and node and link topology. As stated in 5.2.3/G.7715/Y.1706 "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."

Multiple RCs within an RA may transform and then forward information to RCs at different levels. However, in this case, the resulting information at the receiving level shall be self-consistent; this may be achieved using a number of mechanisms.

Since the routing paradigm of RPs may be different between routing levels, information flows between levels shall not be specific to any paradigm (e.g., centralized or link state). An RP using a link state protocol shall support the passing of reachability and may support the passing of topology information to and from its adjacent levels.

## **8.1.1 Communication between adjacent routing levels**

In order to support multilevel hierarchical routing, routing functions at different levels communicate and information elements are exchanged.

### **8.1.1.1 Type of information exchanged**

As specified in ITU-T Rec. G.7715/Y.1706, the type of information flowing upward (i.e., level N to level N+1) and the information flowing downward (i.e., level N+1 to level N) are used for similar purposes, namely, the exchange of reachability information, and may include summarized topology information to allow routing across multiple RAs. It should be noted that summarization of topology information may have an impact on the accuracy of routing and may require additional path calculation outside of the local RP. Clauses 8.1.1.2 and 8.1.1.3 further describe potential methods of upward and downward communication.

### **8.1.1.2 Level N+1 visibility to level N reachability and topology**

Two approaches are described here for upward communication. Other approaches are not precluded and are for further study.

In the first approach, the level N+1 routing function is statically configured with the endpoints located in level N and some abstracted topology of nodes and links. The reachability information may be represented by an address prefix to facilitate scalability, or it may be an actual list of the endpoints in the RA.

In the second approach, the level N+1 RC listens to the routing protocol exchange occurring in each contained level N RA and retrieves the endpoints being announced by the level N routing instance(s) and the full level N topology. This information may be summarized into one or more address prefixes and an abstracted topology in order to facilitate scalability.

Either approach may be used; however, the dynamic approach allows greater flexibility to advertise changes in the RAs at level N.

### **8.1.1.3 Level N visibility to level N+1 reachability and topology**

Two approaches are described here for downward communication. Other approaches are not precluded and are for further study.

In the first approach, the RP in the containing RA at level N+1 provides the level N RP with the reachability and topology information visible at level N+1. The information visible at level N+1 includes the information visible at consecutive upper levels. This information may then be used by the level N RP to compute a path that leaves the RA.

In the second approach, recursive requests are made from the level N RP to the level N+1 RP upward towards the root of the routing hierarchy. The result of each request is analysed by the requesting RP to determine the exit point utilized by the level N+1 RP. The RP will then update the path including the path computed through the level N RA, and return it to the requester. This approach is loop-free as the routing hierarchy defined in ITU-T Rec. G.8080/Y.1304 has strict containment, preventing a contained RA from containing a containing RA.

Either approach may be used for an RC in an RA at level N to develop paths to points outside of the RA.

### **8.1.1.4 Interactions between upward and downward communication**

When both the upward and downward communication interfaces contain endpoint reachability information, a feedback loop could potentially be created. Consequently, the link state routing protocol shall include a method to:

- prevent information propagated from a level N+1 RA into the level N RA to be re-introduced into the level N+1 RA; and

- prevent information propagated from a level N-1 RA into the level N RA to be re-introduced into the level N-1 RA.

In order to prevent such potential loops, there is a requirement for the routing protocol to differentiate between routing information generated within the level of the receiving RC and information that has been received from higher or lower levels, even when this is forwarded by another RC at the same level.

#### **8.1.1.5 Method of communication**

Two approaches exist for communication between level N and level N+1. The first approach places an instance of a level N routing function and an instance of a level N+1 routing function in the same system. The communications interface is within a single system and is thus not an open interface subject to standardization.

The second approach places the level N routing function on a separate system from the level N+1 routing function. In this case, the relationship between the systems containing the routing functions for different levels needs to be established by, e.g., configuration or automated discovery, and a protocol shall be used between the systems. The mechanisms and protocol associated with this interface is not within the scope of this Recommendation.

### **8.2 LRM to RC communications**

One of the responsibilities of the LRM is to provide the RC with information regarding the type and availability of resources on a link, and any changes to those resources.

This requires the following basic functions between the LRM and the RC:

- The RC populates its RDB with information it receives from the LRM.
- The LRM informs the RC of link changes.
- The RC advertises this information to adjacent RCs. A method shall exist to ensure that other RCs can distinguish between up-to-date and out-of-date information.

NOTE – The LRM should utilize procedures to prevent RC overloading. Advertisements between RCs should utilize procedures to avoid RC overloading.

### **8.3 Configuring the hierarchy and information flow**

The RC shall support static (e.g., operator-assisted) configuration and it may support automated configuration of the information describing its relationship to parent and child RPs within the hierarchical routing structure (including RA ID and RC ID). When applied recursively, the whole hierarchy is thus configured.

### **8.4 Configuring RC adjacencies**

The RC shall support static (e.g., operator-assisted) configuration and may provide automated configuration of the information describing its control adjacencies to other RCs within an RA.

The protocol should support all the types of adjacencies described in clause 9/G.7715/Y.1706.

## **9 Routing attributes**

### **9.1 Principles**

ITU-T Rec. G.8080/Y.1304 describes the set of control plane components that are used to manipulate transport network resources in order to provide the functionality of setting up, maintaining and releasing network layer connections. Routing for transport networks is performed on a per-layer basis, where the routing paradigms may differ among layers. Not all transport equipments support the same set of transport layers, nor the same degree of connection flexibility at

any given layer. A server layer trail may support various clients, involving different adaptation functions. Additionally, as described in ITU-T Rec. G.8080/Y.1304 Amendment 1, transport equipment may support variable adaptation functionality, whereby a single server layer trail dynamically supports different multiplexing structures. As a result, path computation and routing are impacted by layer-specific, layer-independent, and client/server adaptation information elements. The routing protocol must be applicable to any transport layer network (e.g., G.803, G.872) and the representation of routing attributes should not preclude their applicability to other transport network layers, existing or future.

## **9.2 Taxonomy of attributes**

Following the above architectural principles, attributes can be organized according to the following categories:

- Node-related or link-related.
- Provisioned, negotiated or automatically configured.
- Inherited from the server layer or specified by layer.

Configuration of link attributes can be statically or automatically performed for each transport network layer. Although configuration can be performed separately for each layer, this may lead to unnecessary repetition. Hence, the inheritance property of attributes can also be used to optimize the configuration process.

SNPP links are configured by the operator through grouping of SNP link connections. Grouping may be based on different link attributes (e.g., diversity support, link weight, etc). Configuration operations (e.g., configuration of link attributes or of SNPPs) must be designed such that the operation and any subsequent routing protocol advertisement does not affect existing transport plane connections.

Two RAs may be linked by one or more SNPP links. Multiple SNPP links may be required when SNP link connections are not equivalent for routing purposes with respect to:

- the RAs to which they are attached; or
- the containing RA; or
- when smaller groupings are required for administrative purposes.

## **9.3 Common advertisement information**

All advertisements contain a common set of administrative information elements. These elements are:

- RA ID of which the advertisement is bounded.
- RC ID of the entity generating the advertisement.
- Information to uniquely identify advertisements.
- Information to determine whether an advertisement has been updated.
- Information to indicate when an advertisement comes from a different level (which represents information external to the RA).

## **9.4 Node attributes**

All nodes in the graph representation of a network belong to an RA, hence, the RA ID is an attribute of all nodes. Given that no distinction is made between abstract nodes and those that cannot be decomposed any further, as described in clause 5, the same attributes are used for their advertisement.

The following node attributes are defined:

- Node ID;
- Reachability Information: Reachability information describes the set of endpoints that are reachable by the associated node. It may be advertised either as a set of UNI transport resource addresses/address prefixes, or a set of associated SNPP IDs/SNPP ID prefixes, the selection of which must be consistent within the applicable scope.

In Table 1, capability refers to level of support required in the realization of a link state routing protocol, whereas usage refers to degree of operational and implementation flexibility.

**Table 1/G.7715.1/Y.1706.1 – Node attributes**

Node attribute	Capability	Usage
Node ID	Mandatory	Mandatory
Reachability	Mandatory	Optional

## 9.5 Link attributes

The following link attributes are defined:

- Local SNPP name;
- Remote SNPP name;
- Layer-specific characteristics (refer to 9.5.1).

In Table 2, capability refers to level of support required in the realization of a link state routing protocol, whereas usage refers to degree of operational and implementation flexibility.

**Table 2/G.7715.1/Y.1706.1 – Link attributes**

Link attribute	Capability	Usage
Local SNPP name	Mandatory	Mandatory
Remote SNPP name	Mandatory	Mandatory
Layer-specific characteristics	(refer to 9.5.1)	

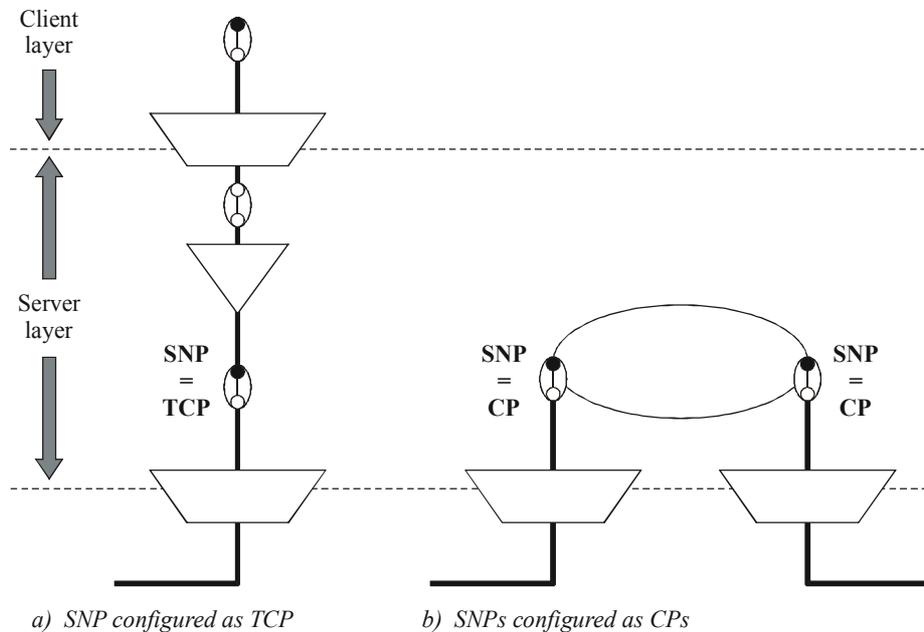
NOTE – When the remote end of a link is located outside of the RA, usage of the remote SNPP name is optional.

### 9.5.1 Layer-specific characteristics

- Signal type: This attribute identifies the characteristic information of the layer network.
- Link weight: This attribute represents a vector of one or more metrics, each of which indicates the relative desirability of a particular link over another during path selection.
- Resource class: This attribute corresponds to a set of administrative groups assigned by the operator to this link. A link may belong to zero, one or more administrative groups.
- Local connection type: This attribute identifies whether the local SNP represents a TCP, CP, or can be flexibly configured as either a TCP or a CP. Figure 4 illustrates this attribute.
- Link capacity: This attribute provides the sum of the available and potential link connections for a particular network transport layer. Other capacity information may be useful and is not precluded.
- Link availability: This attribute represents a vector of one or more availability factors for the link, or link end. Availability may be represented in different ways between domains and within domains. Within domains, it may be used to represent a survivability capability

of the link or link end. In addition, the availability factor may be used to represent a node survivability characteristic.

- Diversity support: This attribute represents diversity information with respect to links, nodes and Shared Risk Groups (SRGs) that may be used during path computation.
- Local client adaptations supported: This attribute represents the set of client layer adaptations supported by the TCP associated with the local SNPP. This is only applicable when the local SNP represents a TCP or can be flexibly configured as either a TCP or CP.



G.7715.1/Y.1706.1\_F04

**Figure 4/G.7715.1/Y.1706.1 – SNP realizations**

NOTE – Separate layer advertisements of layer-specific attributes may be chosen, however, this may lead to unnecessary duplication. This can be avoided using the local adaptation information and the inheritance property described in 9.2.

The interlayer information advertisement is achieved through the coordination of the LRMs responsible for SNPPs at each layer.

In Table 3, capability refers to level of support required in the realization of a link state routing protocol, whereas usage refers to degree of operational and implementation flexibility.

**Table 3/G.7715.1/Y.1706.1 – Layer-specific characteristics**

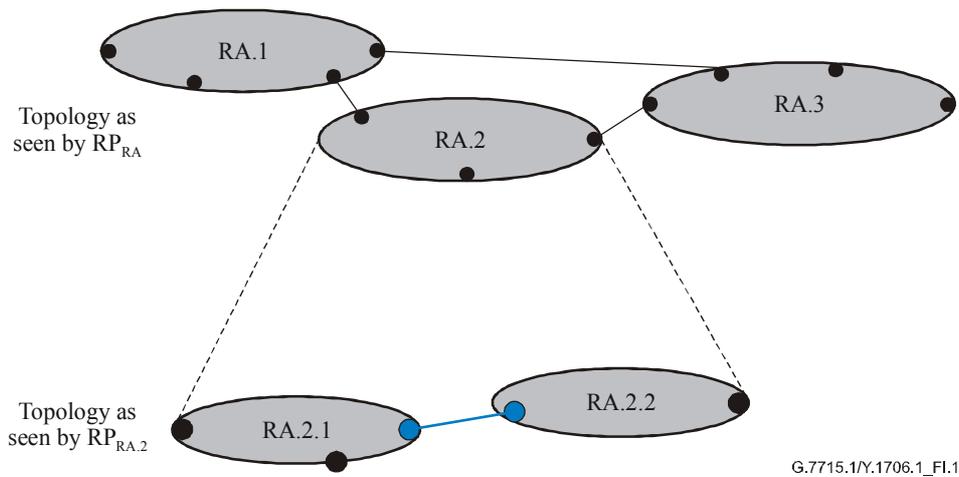
Layer-specific characteristics	Capability	Usage
Signal type	Mandatory	Optional
Link weight	Mandatory	Optional
Resource class	Mandatory	Optional
Local connection type	Mandatory	Optional
Link capacity	Mandatory	Optional
Link availability	Optional	Optional
Diversity support	Optional	Optional
Local client adaptations supported	Optional	Optional

## Appendix I

### Use of SNPP aliases in multilevel routing hierarchies

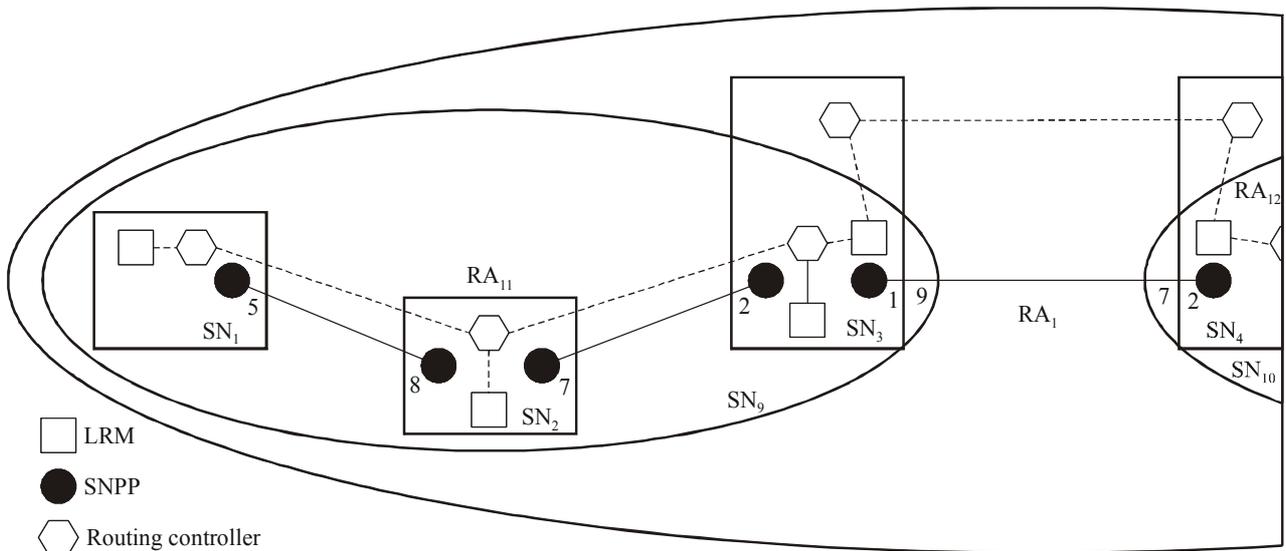
This appendix describes an illustrative scenario, with a multilevel routing hierarchy using a separate SNPP name space per level, to illustrate that SNPP aliases are useful to routing.

Routing areas may be hierarchically contained, with a separate routing performer associated with each routing area in the hierarchy. Since the RP for this RA only has visibility of the topology of its RA, it has no specific knowledge of the topology of RAs that contain it, or any of the RAs it contains. The RP may represent the contained RA as an abstract node along with the SNPP links that connect the contained RA to other RAs. Figure I.1 (also Figure 7/G.7715/Y.1706) illustrates such a topology.



**Figure I.1/G.7715.1/Y.1706.1 – Topology views as seen by RP associated with hierarchical RAs (Figure 7/G.7715/Y.1706)**

The definition of an RA in ITU-T Rec. G.8080/Y.1304 requires that links be wholly contained within an RA. Consequently, a link does not exist in any RA other than the lowest RA that contains both endpoints of a link. An RC for an RA containing the link may not be collocated with the LRM responsible for the link. This is illustrated in Figure I.2.



G.7715.1/Y.1706.1\_F1.2

**Figure I.2/G.7715.1/Y.1706.1 – Use of SNPP aliases in contained areas**

Two separate SNPP names exist for the link end in SN<sub>3</sub> that is connected to SN<sub>4</sub>:

$$RA \leq RA_{11}, RA_{11} > SN = SN_3 \quad LC = 1 \quad (\text{in the } RA_{11} \text{ context})$$

and

$$RA \leq RA_1 > SN = SN_9 \quad LC = 9 \quad (\text{in the } RA_1 \text{ context})$$

Since RA<sub>1</sub> is the lowest RA that contains both link endpoints, the link only exists in RA<sub>1</sub>. Use of the link will be as a result of the RP for RA<sub>1</sub> providing the RA<sub>1</sub> name for the SNPP.

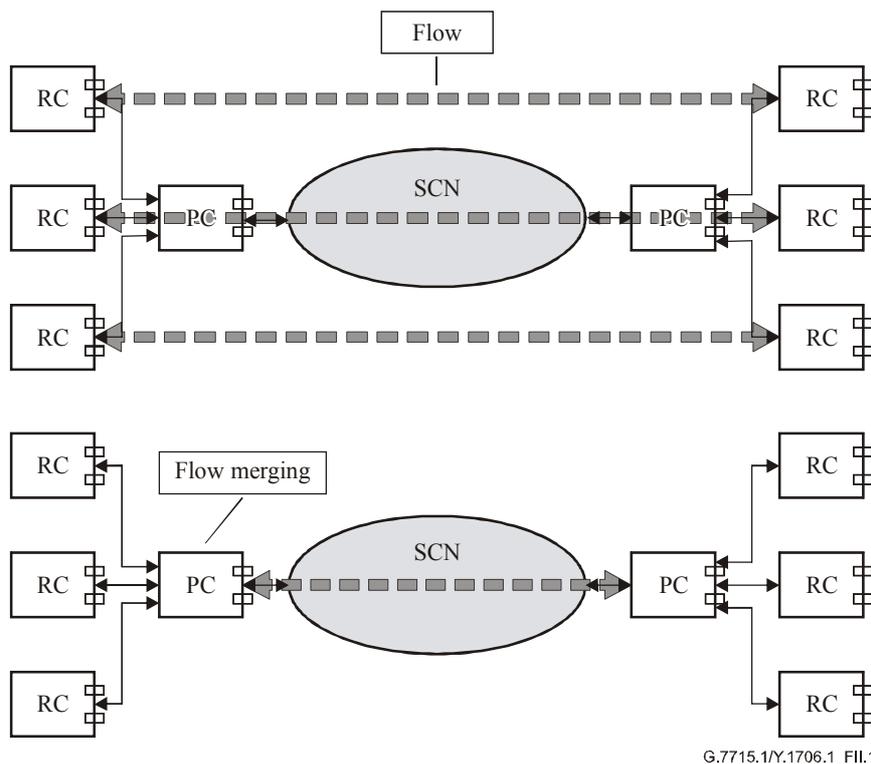
When hierarchical path computation is performed from a point in RA<sub>11</sub> to a point outside of RA<sub>11</sub>, the RA<sub>11</sub> RP will be asked to develop a path to an egress node using the RA<sub>1</sub> name that it does not recognize. Consequently, the RA<sub>1</sub> name will need to be translated into an RA<sub>11</sub> name. The approach for this translation is protocol-dependent and, therefore, outside the scope of this Recommendation.

## Appendix II

### Multiple RC communications over PCs

This appendix provides illustrative examples of different methods of carrying multiple layer information between two systems. It does not mandate a particular method of implementation for link state routing protocols.

From the perspective of ITU-T Rec. G.8080/Y.1304, each layer network has its own instance of a control plane. When multiple layer networks exist, RCs at each layer exchange information. Information from a set of RCs can be exchanged by a single PC. Two examples of this information flow are described below. The first one employs a protocol controller which performs multiplexing of RC-to-RC communication for each pair of RCs. The other one is to have a multiplexing of the RC-to-RC information by the PC-to-PC communication. This is shown in Figure II.1:



G.7715.1/Y.1706.1\_FII.1

**Figure II.1/G.7715.1/Y.1706.1 – Protocol controller multiplexing pairs of RC communication (upper figure) and protocol controller multiplexing RC-to-RC information together (lower figure)**

### Appendix III

#### Example scenarios for RA segmentation, aggregation and evolution of the network hierarchy

This appendix provides illustrative examples of operations that may be performed on an RA. It does not mandate a particular method of implementation for link state routing protocols.

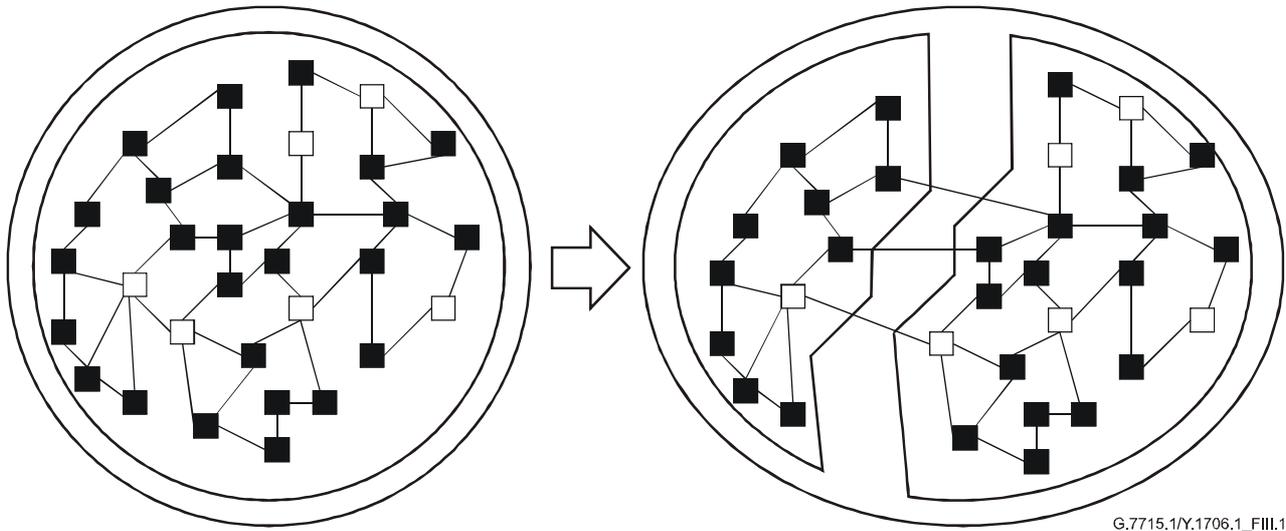
A number of operations may be performed on an RA, including the following:

- Segmentation of one routing area into two separate RAs;
- Aggregation of two RAs into one RA;
- Renaming of RAs;
- Insertion of an RA into the hierarchy;
- Deletion of an RA from the hierarchy.

There are many different scenarios that may require these operations. The following clauses describe them.

### III.1 Segregation – Insertion of a node in an RA that is "full"

RAs may be restricted in the number of nodes they can support. This constraint may be due to vendor implementation limitations, or network operator practices.



**Figure III.1/G.7715.1/Y.1706.1 – Split of an RA with 32 nodes**

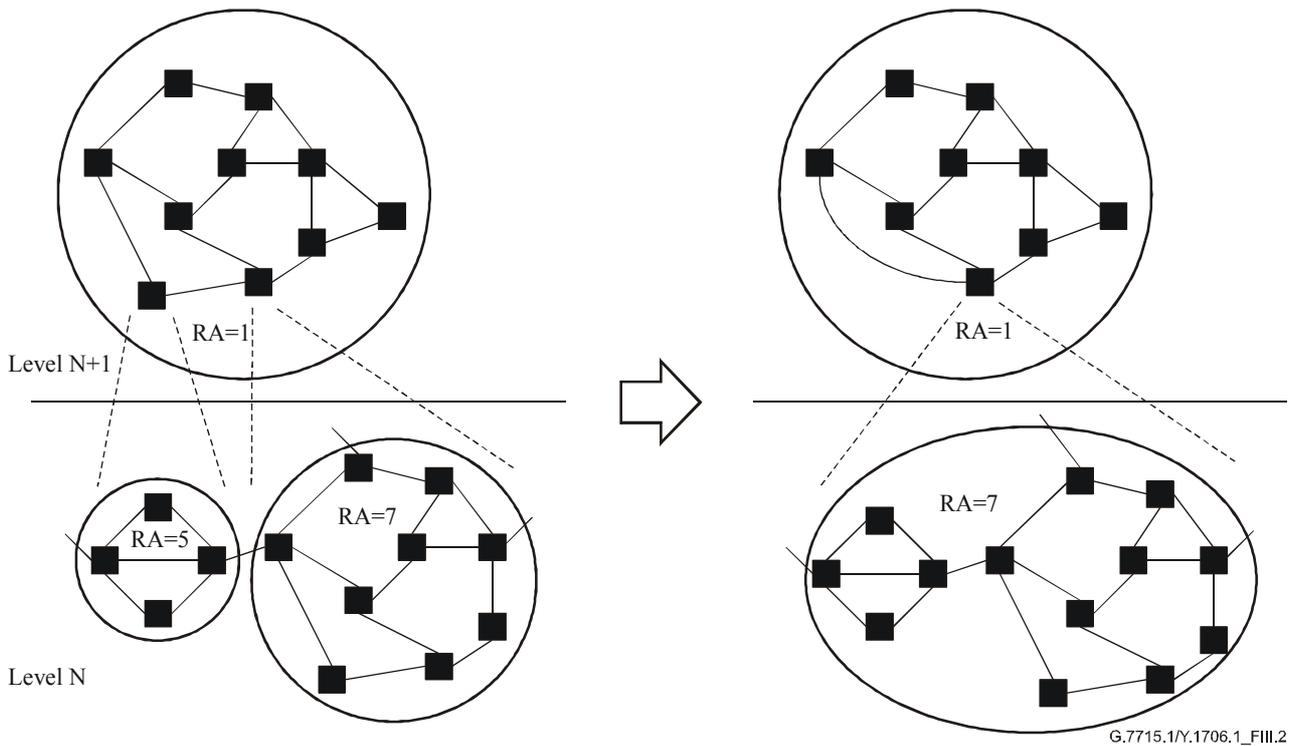
The RA shown in Figure III.1 contains 32 nodes, some of which (shown in white) are unable to handle RAs larger than 32 nodes in size. This may be, for example, due to the availability of memory or CPU to handle the routing protocol. Adding a node to this RA will likely cause a service interruption for the RCs on the white nodes.

After splitting, the resulting two RAs have 11 nodes in one RA, and 21 nodes in the other. Therefore, the new node may be added to either RA without exceeding the limitations of the RCs on the white nodes.

### III.2 Merging two RAs

Two RAs may be merged into a single RA. This may be due to an interest in flattening a topological hierarchy, or due to business activities such as a merger or an acquisition.

As an example, Figure III.2 shows two RAs that are to merge. Both of these RAs service the same geographical region, and are a part of the same administrative domain, making it unnecessary for them to stay in separate RAs.

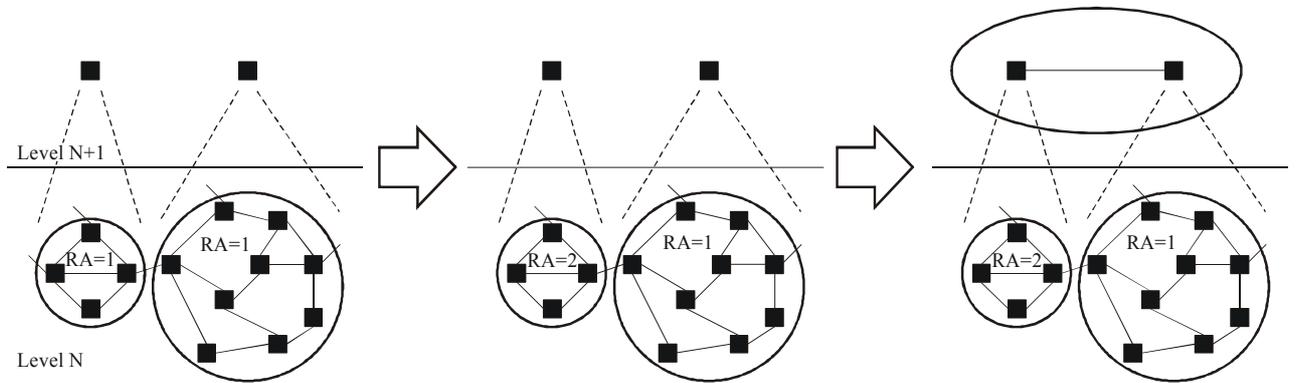


G.7715.1/Y.1706.1\_FIII.2

Figure III.2/G.7715.1/Y.1706.1 – Merging of two RAs

### III.3 Renaming an RA

It is only possible to connect two RAs after verifying that their corresponding ID is unique across both networks. Any RA IDs that are used in both RAs will need at least one of the RA IDs to be renamed. After renaming, the RAs can be interconnected. This progression is shown in Figure III.3.



G.7715.1/Y.1706.1\_FIII.3

Figure III.3/G.7715.1/Y.1706.1 – Renaming RAs

### III.4 Insertion – Different routing protocol added to a portion of a network

One example of insertion involves the addition of a different routing protocol to a portion of a network. For instance, a network operator may decide to deploy equipment in a portion of their network which is unable to participate in the protocol communications currently used in that portion. Consequently, an RA that contains RAs for each protocol type can be used to summarize reachability. This new RA is inserted above the existing RA, as shown in Figure III.4.

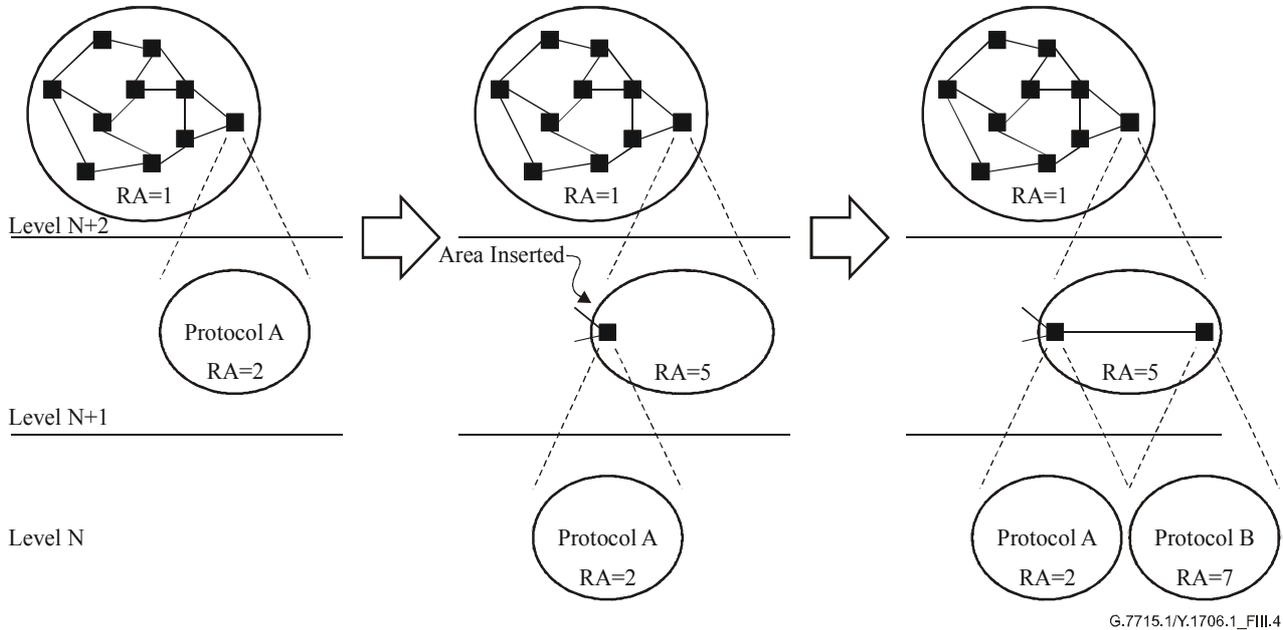


Figure III.4/G.7715.1/Y.1706.1 – Adding a different protocol to a portion of the network

ITU-T Y-SERIES RECOMMENDATIONS  
**GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT  
GENERATION NETWORKS**

<b>GLOBAL INFORMATION INFRASTRUCTURE</b>	
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
<b>INTERNET PROTOCOL ASPECTS</b>	
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
<b>Operation, administration and maintenance</b>	<b>Y.1700–Y.1799</b>
Charging	Y.1800–Y.1899
<b>NEXT GENERATION NETWORKS</b>	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899

*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
<b>Series G</b>	<b>Transmission systems and media, digital systems and networks</b>
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
<b>Series Y</b>	<b>Global information infrastructure, Internet protocol aspects and Next Generation Networks</b>
Series Z	Languages and general software aspects for telecommunication systems