

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.7701

(11/2016)

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

Data over Transport – Generic aspects – Transport
network control aspects

Common control aspects

Recommendation ITU-T G.7701

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
TRANSMISSION MEDIA AND OPTICAL SYSTEMS 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
MULTIMEDIA QUALITY OF SERVICE AND PERFORMANCE – GENERIC AND USER-RELATED ASPECTS	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DATA OVER TRANSPORT – GENERIC ASPECTS	G.7000–G.7999
General	G.7000–G.7099
Transport network control aspects	G.7700–G.7799
PACKET OVER TRANSPORT ASPECTS	G.8000–G.8999
ACCESS NETWORKS	G.9000–G.9999

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

Recommendation ITU-T G.7701

Common control aspects

Summary

Recommendation ITU-T G.7701 describes concepts that are common to both software defined networking (SDN) controller and automatically switched optical network (ASON) control approaches, including common aspects of the interaction between the control functions, management functions and transport resources.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T G.7701	2016-11-13	15	11.1002/1000/13090

Keywords

ASON, control components, management-control continuum, transport SDN.

* To access the Recommendation, type the URL <http://handle.itu.int/> in the address field of your web browser, followed by the Recommendation's unique ID. For example, <http://handle.itu.int/11.1002/1000/11830-en>.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). 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, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2017

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

Table of Contents

		Page
1	Scope.....	1
2	References.....	1
3	Definitions	2
	3.1 Terms defined elsewhere	2
	3.2 Terms defined in this Recommendation.....	2
4	Abbreviations and acronyms	4
5	Conventions	6
6	Overview.....	6
7	Transport resources and their representation	7
	7.1 Transport functional architecture	7
	7.2 Domains.....	8
	7.3 Control view of transport resources for connection management.....	8
	7.4 Virtualization.....	10
8	Control component approach.....	11
	8.1 Notation	11
	8.2 Policy.....	13
	8.3 Common components	15
9	Common control communications.....	35
10	Common management aspects of common control components	36
	10.1 Control component relationships.....	36
	10.2 Control function management requirements	37
11	Identifiers.....	37
	11.1 Resources in the transport network	37
	11.2 Control view of transport resources	37
	11.3 Control components.....	38
	11.4 Control artefacts	38
	11.5 Reference points	38
	11.6 Control communications network	38
	Annex A – Configuration of OTN digital and media layers.....	39
	Bibliography.....	42

Recommendation ITU-T G.7701

Common control aspects

1 Scope

This Recommendation describes concepts that are common to both software defined networking (SDN) controller and automatically switched optical network (ASON) control approaches. This encompasses common aspects of:

- transport resources and their representation;
- control components used to describe architecture for application of SDN and ASON to the control of transport resources;
- control communication aspects;
- common management aspects of common control components;
- identifiers with respect to naming and addressing, which involves separation of identity from location.

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 G.698.1] Recommendation ITU-T G.698.1 (2009), *Multichannel DWDM applications with single-channel optical interfaces*.
- [ITU-T G.698.2] Recommendation ITU-T G.698.2 (2009), *Amplified multichannel dense wavelength division multiplexing applications with single channel optical interfaces*.
- [ITU-T G.800] Recommendation ITU-T G.800 (2016), *Unified functional architecture of transport networks*.
- [ITU-T G.805] Recommendation ITU-T G.805 (2000), *Generic functional architecture of transport networks*.
- [ITU-T G.872] Recommendation ITU-T G.872 (2012), *Architecture of optical transport networks*.
- [ITU-T G.7711] Recommendation ITU-T G.7711/Y.1702 (2015), *Generic protocol-neutral information model for transport resources*.
- [ITU-T G.7712] Recommendation ITU-T G.7712/Y.1703 (2010), *Architecture and specification of data communication network*.
- [ITU-T G.7714.1] Recommendation ITU-T G.7714.1/Y.1705.1 (2015), *Protocol for automatic discovery in transport networks*.
- [ITU-T G.8080] Recommendation ITU-T G.8080/Y.1304 (2012), *Architecture for the automatically switched optical network*.
- [ITU-T M.3100] Recommendation ITU-T M.3100 (2005), *Generic network information model*.

- [ITU-T Q.2982] Recommendation ITU-T Q.2982 (1999), *Broadband integrated services digital network (B-ISDN) – Digital subscriber signalling system No. 2 (DSS2) – Q.2931-based separated call control protocol*.
- [ITU-T X.731] Recommendation ITU-T X.731 (1992), *Information technology – Open Systems Interconnection – Systems management: State*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

- 3.1.1 **access point**: [ITU-T G.800].
- 3.1.2 **adaptation**: [ITU-T G.800].
- 3.1.3 **characteristic information (CI)**: [ITU-T G.800].
- 3.1.4 **connection**: [ITU-T G.800].
- 3.1.5 **data communications network**: [ITU-T G.7712].
- 3.1.6 **forwarding point (FP)**: [ITU-T G.800].
- 3.1.7 **layer network**: [ITU-T G.805].
- 3.1.8 **link**: [ITU-T G.805].
- 3.1.9 **link connection**: [ITU-T G.805].
- 3.1.10 **logical termination point**: [ITU-T G.7711].
- 3.1.11 **software defined networking**: [b-ITU-T Y.3300].
- 3.1.12 **subnetwork**: [ITU-T G.805].
- 3.1.13 **subnetwork connection**: [ITU-T G.805].
- 3.1.14 **trail**: [ITU-T G.805].
- 3.1.15 **transitional link**: [ITU-T G.800].

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

- 3.2.1 **address**: A string of symbols that is valid regardless of the location of the source but changes if the destination moves. An address is used for the purpose of routing. Source and destination addresses must be globally unique.
- 3.2.2 **administrative domain**: A type of domain that belongs to a single entity such as a network operator, a service provider or an end-user.
- 3.2.3 **boundary resource identifier (BRI)**: Used at a domain boundary to identify both the transport resources, and the interface(s) between control components, at that boundary.
- 3.2.4 **call**: An association between two or more users and one or more domains that support an instance of a service through one or more domains. Within domains, the association is supported by network entities that contain call state. Between a user and a network call control entity and between network call control entities, there are call segments. The call consists of a set of concatenated call segments.
- 3.2.5 **call admission control**: A policy function invoked by an originating role in a network and may involve cooperation with the terminating role in the network.

3.2.6 call controller: Calls are controlled by means of call controllers. There are two types of call controller components.

3.2.7 calling/called party call controller: The control component associated with an end of a call and may be co-located with end systems or located remotely and acts as a proxy on behalf of end systems. This controller acts in one or both of two roles, one to support the calling party and the other to support the called party.

3.2.8 call segment: An association between two call control entities (as per [ITU-T Q.2982]). Each call segment has zero or more associated connections. Call segments between network call control entities have zero or more supporting calls.

3.2.9 component: An abstract representation of a functional entity. In this Recommendation, components do not represent instances of implementation code. They are used to construct scenarios to explain the operation of the architecture.

3.2.10 component interface: Interfaces that represent the logical relationships between control components and are defined by the information flow between these entities.

NOTE – Such a relationship allows for choice in selection of potentially exposable interfaces in support of different equipment implementations and network architectures.

3.2.11 connection admission control: A process that determines if there are sufficient resources to admit a connection (or re-negotiates resources during a call).

3.2.12 connection controller (CC): The control component responsible for coordination among the link resource manager, the routing controller and both peer and subordinate connection controller components for the purpose of the management and supervision of connection setups, releases and the modification of connection parameters for existing connections.

3.2.13 control domain: A type of domain that represents the scope of control over a collection of resources allocated to that domain as determined by a set of functions from the management-control continuum.

3.2.14 domain: Represents a collection of resources that are grouped for a particular purpose. Domains of the same type do not overlap among themselves, but a hierarchical containment relationship of domains is possible.

NOTE – In general, resources may be considered to encompass compute, networking, and storage.

3.2.15 network call controller (NCC): A control component providing three roles, one for support of the calling party, another to support the called party and a third to support calls across domain boundaries.

3.2.16 policy: The set of rules selected from among alternatives and in the light of given conditions to guide and, usually, to determine present and future decisions.

3.2.17 recovery domain: A type of control domain whose purpose is assuring the reliable transfer of information across the resources allocated to that domain.

NOTE – Approaches to support reliable transfer of information may include, e.g., protection, distributed or centralized restoration, etc. A pre-condition for establishing a recovery domain is that there are sufficient network resources in the domain to maintain a specified level of reliability for the selected approach(es). In multi-layer transport networks, a recovery domain in a client layer network must fully contain any recovery domains in the server layer networks that it uses.

3.2.18 resource database (RDB): A logical entity that holds (makes available) information that components use for their operations. Related components share the same RDB.

3.2.19 route: A sequence of the artefacts representing the transport resources (i.e., SNPs, SNPPs, and routing areas) that support, or are intended to support, a connection.

3.2.20 routing area (RA): An instance of a routing domain, and is defined by a set of subnetworks, the SNPP links that interconnect them, and the SNPPs representing the ends of the SNPP links exiting that routing area. A routing area may contain smaller routing areas interconnected by SNPP links. The limit of subdivision results in a routing area that contains a subnetwork.

3.2.21 routing controller (RC): A control component with the roles to:

- respond to requests for path (route) information needed to set up connections. This information can range from end-to-end path details to a next hop. The route can be computed by one or more cooperating RCs;
- respond to requests for topology (SNPs and their abstractions) information for management-control continuum purposes.

3.2.22 routing level: A relationship between a routing area (RA) and a containing RA or contained RAs. The containment hierarchy of routing areas creates routing levels.

3.2.23 routing domain: A type of control domain whose purpose is routing across the collection of resources allocated to that domain.

3.2.24 service level agreement: A contract between two parties such as a service provider and a customer. It defines the services available to the customer, and the grade of service of those services as offered to the customer. It also usually describes the service guarantee and potential penalties in case of service degradation or failure.

3.2.25 subnetwork point (SNP): A control abstraction that represents an actual or potential underlying forwarding point (FP) or forwarding end point (FwEP) (see [ITU-T G.800]). Several SNPs (in different VNs) may represent the same FP or FwEP.

3.2.26 subnetwork point pool (SNPP): A set of SNPs that are grouped together.

NOTE – An SNPP usually represents a link end. An SNPP may be further subdivided (sub-structured) into smaller pools.

3.2.27 virtual network (VN): A designated subset of abstracted network resources.

NOTE – The network resources in the VN may be at different levels of abstraction and may correspondingly use identifiers from different name spaces.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

ASON	Automatically Switched Optical Network
BRI	Boundary Resource Identifier
CallC	Call Controller
CAC	Connection Admission Control
CC	Connection Controller
CCC	Calling/called party Call Controller
CCI	Connection Controller Interface
CCN	Control Communications Network
CI	Characteristic Information
CIR	Committed Information Rate
cFP	Client Forwarding Point
CTP	Connection Termination Point

DA	Discovery Agent
DCN	Data Communication Network
DM	Discovery Message
EMS	Element Management System
FP	Forwarding Point
FwEP	Forwarding End Point
IM	Information Model
LRM	Link Resource Manager
LP	Layer Protocol
LTP	Logical Termination Point
MCC	Management-Control Continuum
MPLS	Multiprotocol Label Switching
NCC	Network Call Controller
NE	Network Element
NMS	Network Management System
OAM	Operations, Administration and Maintenance
OCN	Overhead Communications Network
OSS	Operations Support System
OTN	Optical Transport Network
OTSi	Optical Tributary Signal
OTSiA	Optical Tributary Signal Assembly
PC	Protocol Controller
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PIR	Peak Information Rate
RA	Routing Area
RC	Routing Controller
RDB	Resource Database
SDN	Software Defined Networking
sFP	Server Forwarding Point
SNC	Subnetwork Connection
SNP	Subnetwork Point
SNPP	Subnetwork Point Pool
TAP	Termination and Adaptation Performer
TTP	Trail Termination Point
UML	Unified Modelling Language
UNI	User-Network Interface (reference point)

5 Conventions

This Recommendation uses the diagrammatic conventions defined in [ITU-T G.800] to describe the transport resources.

6 Overview

Considering connection management, network management, ASON control, and SDN controllers all perform the same operations on the transport resources. The distinctions among these approaches are based upon variations that may exist in the degree of automation (high human intervention to none), distribution of their implementation (centralized versus fully or partially distributed), and exposure of interfaces (closed versus open). However, the transport resources themselves make no distinctions based upon how the connection management function is performed but simply respond to the functions operating on them and report changes (via notifications).

The concept whereby management and control functions are considered to be a continuum is known as the management-control continuum (MCC); i.e., it expresses the view that management and control functions are essentially the same and, thus, they can be grouped into one set of MCC functions.

In the MCC, interfaces between functions may be exposed when resources are placed into adjacent domains, when an automated interface to users is desired, when a connection management function operates on a multi-vendor set of resources, as well as other factors.

Figure 6-1 illustrates the MCC concept and its relationship with transport resources, whereby management and control functions operate on transport resources and receive state information about resources. Transport resource management functional areas are identified in [ITU-T M.3010] as: performance management, fault management, configuration management, accounting management and security management.

The management and control functions themselves require management, depicted as the administration role functions in Figure 6-1. SDN controller management provides the ability to configure SDN controllers. ASON control management provides for management of control components such as the routing controller (RC) component. For network management, this role provides the ability for configuration and management of operations support system (OSS)/network management system (NMS)/element management system (EMS).

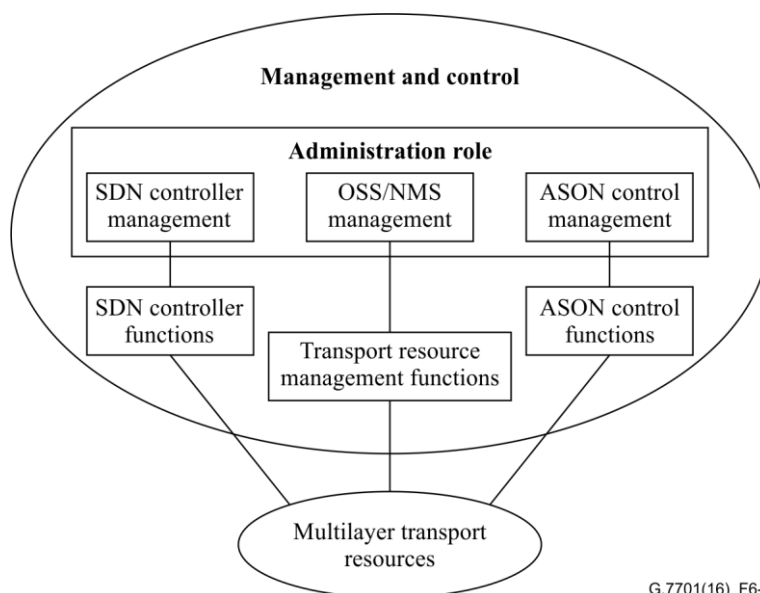


Figure 6-1 – Management-control continuum

Not shown is the data communications network (DCN), which provides the communication paths to carry, e.g., ASON control, SDN controller, and management information communications. The details of the DCN, management functions, and transport resources are specified in other Recommendations and are outside the scope of this Recommendation.

Network resources may be partitioned among those under the authority of management systems, those under the authority of the SDN controllers, and those under the authority of the ASON control. Once a subset of the resources is assigned to a particular MCC instance (e.g., ASON control), other MCC instances (e.g., SDN controller) cannot modify these resources.

This includes transport network resources not currently in use, but reserved for future use (e.g., by network planners). As such, transport resource partitioning is performed by the management function and thus, is outside the scope of this Recommendation.

7 Transport resources and their representation

The transport network is a large, complex network with various components, and an appropriate network model with well-defined, technology agnostic, functional entities is essential for its design, control, and management. The transport network can be described by defining the associations between points in the network. The resultant logical network topology allows the separation between logical connections and the physical routes and resources used.

7.1 Transport functional architecture

The functional architecture of the transport network describes the way that the transport resources are used to perform the basic transport functions in a manner that makes no reference to the control and management of those functions. As described in [ITU-T G.805] and [ITU-T G.800], in order to simplify the description, the functional architecture utilizes the concepts of layering and partitioning within each layer network in a manner that allows a high degree of transport network recursion.

- a layer network describes the generation, transport and termination of a particular characteristic information (CI). Layering enables decomposition of a transport network into a number of independent transport layer networks. There is a client/server relationship between each of these layer networks where the client refers to the signal being carried and the server refers to the layer network providing its transport. The client/server paradigm is recursive because any particular server layer could itself be a client of another server layer.
- partitioning is the division of a larger subnetwork into disjoint subnetworks that are interconnected by links. Because the model requires partitions to be nested, partitioning is also recursive.

The components of the transport network architectural model can be divided into three groups: topological components, transport processing functions, and transport entities.

- 1) topological components: Provide a description of a transport network in terms of the relationship between sets of forwarding points (FPs) within a layer network. These encompass, for example, layer network, subnetwork, link (including transitional link), access group.
- 2) transport processing functions: Are used to model the processes implemented in equipment that manipulate the information that is being transferred across the transport network. These include, for example, adaptation, (trail) termination, layer processor, forwarding.
- 3) transport entities: Provide the means to transfer information across the transport network between forwarding points. Transport entities are configured within topological components. These encompass, for example, forwarding relationship, subnetwork transport entity, connection, link connection.

The optical transport network (OTN) has both digital layers and a media layer. The control of optical signals and the media layer must take into account the limitations imposed by the accumulation of analogue impairments within the media. Many of these impairments and the magnitude of their effects are associated with particular technological implementations and the network topology. Annex A provides some clarification on the control of the OTN media layer.

7.2 Domains

ASON control and SDN controller deployments will occur within the context of network operator business practices and the multi-dimensional heterogeneity of transport networks. These business and operational considerations lead to the need for architectural support of, for example, strong abstraction barriers to protect commercial business operating practices, segmenting transport networks into domains according to managerial and/or policy considerations, and inherent transport network heterogeneity (including control and management).

The domain notion referred to is the generalization [ITU-T G.8080] of the [ITU-T G.805] definition of administrative domain and the Internet administrative regions (e.g., autonomous systems) to express differing administrative and/or managerial responsibilities, trust relationships, addressing schemes, infrastructure capabilities, survivability techniques, distributions of control functionality, etc. Domains are established by operator policies and have a range of membership criteria, as exemplified above.

The scope (or boundary) of a domain is defined for a particular purpose, and domains defined for one purpose need not coincide with domains defined for another purpose. Domains that have been defined for the same purpose are restricted in that they do not overlap; however, they may:

- fully contain other domains that have been defined for the same purpose;
- border each other;
- be isolated from each other.

Examples of domains include administrative domains, control domains, routing domains, and recovery domains. Resources may be shared among domains of the same type; an example of this is the shared resources in a virtual network (VN), as described in clause 7.4.

The current assumption is that a control domain must be contained in an administrative domain. Investigation of a peer relationship between control domains that are in different administrative domains is for further study.

7.3 Control view of transport resources for connection management

Connectivity is the basic service provided by a transport network. Connection management functions include: path computation, connection creation, connection modification, connection teardown, and configuration and activation of operations, administration and maintenance (OAM) and survivability mechanisms. These connection management functions could be operated under instructions from an SDN controller, management system, or ASON control.

The underlying transport resources used to support aspects of connection management to the SDN controller and ASON control functions are represented by a set of entities, termed subnetwork point (SNP) and subnetwork point pool (SNPP). The SNP and SNPP entities are organized into routing area (RA), subnetwork and link topological constructs which represent the view of the transport resources as seen by the SDN controller and ASON control components from a connection management perspective. Figure 7-1 illustrates the relationship of the SNP to the transport resources described in [ITU-T G.800] and, the entities that represent these resources from the perspective of network management (as described in [ITU-T M.3100]).

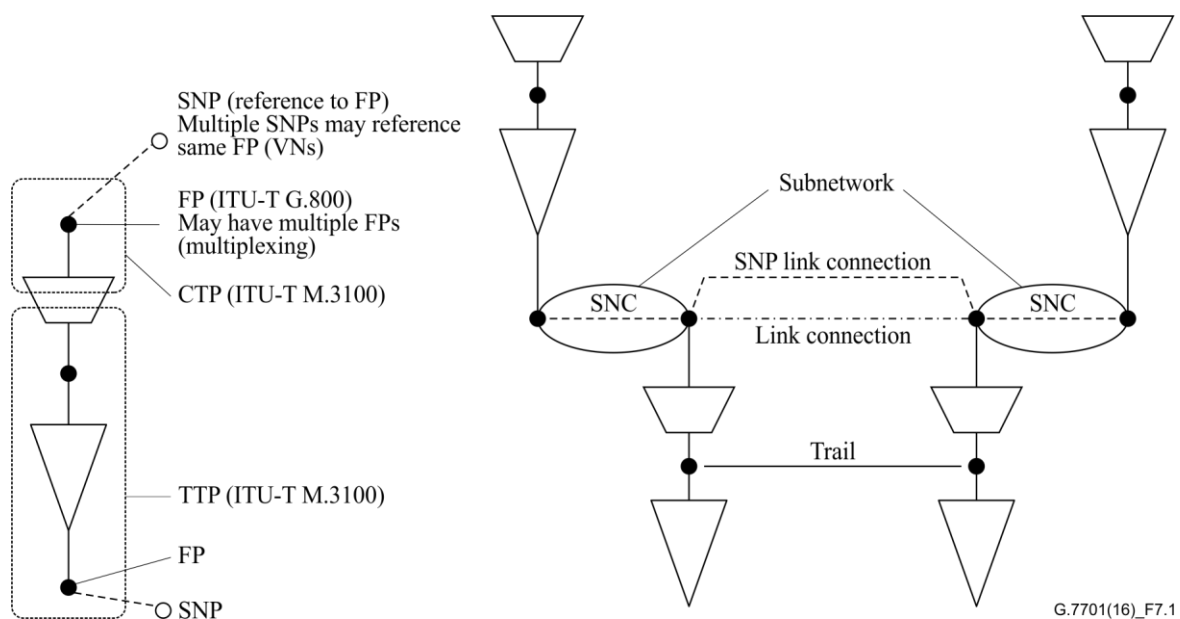


Figure 7-1 – Relationship between architectural entities in the control-management continuum

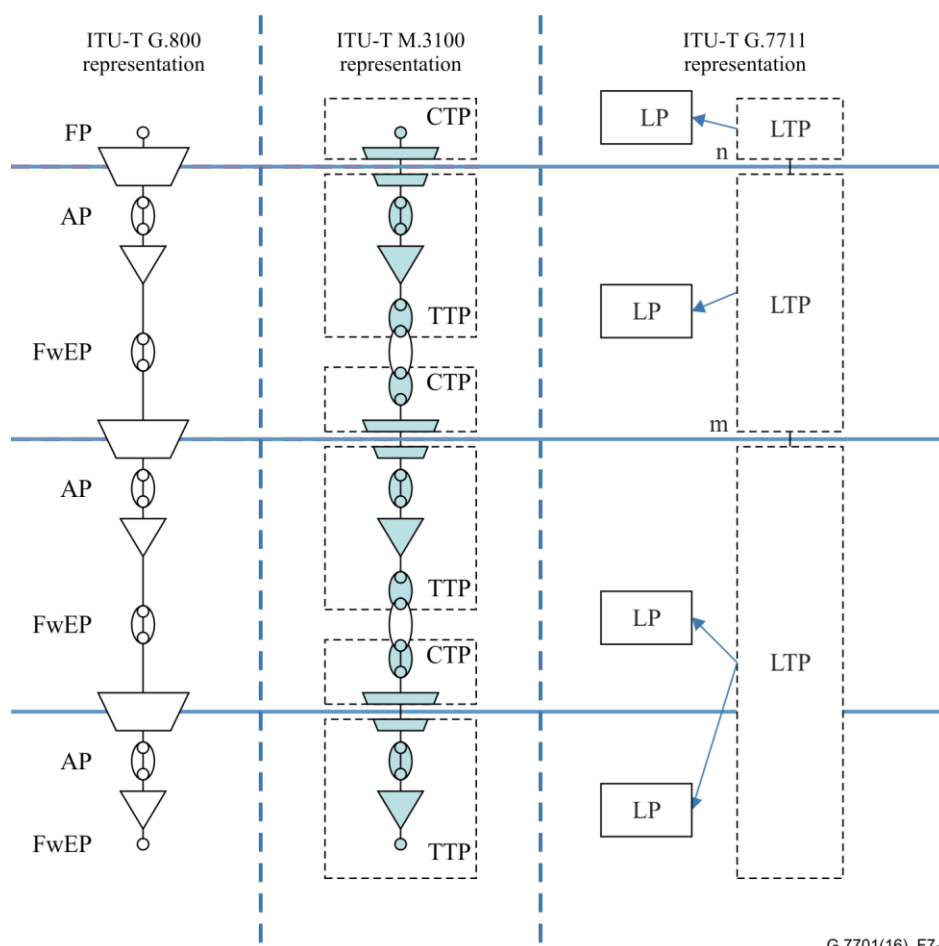
An SNP has a number of relationships with other SNPs:

- a static relationship between two SNPs in different subnetworks. This is referred to as an SNP link connection.
- a dynamic relationship between two (or more in the case of broadcast connections) SNPs at the boundary of the same subnetwork. This is referred to as a subnetwork connection (SNC).

An SNP may also be grouped with other SNPs for the purpose of routing. This is an SNPP and has a strong relationship with link ends; however, this relationship is more flexible than the link end. An SNPP may be further subdivided into smaller pools. One use of this sub-structuring is to describe different degrees of route diversity. For example, all the SNPs in one subnetwork that have a relationship to a similar group on another subnetwork may be grouped into a single SNPP. This SNPP may be further subdivided to represent diverse routes and further subdivided to represent, for example, individual wavelengths.

The association between SNPPs on different subnetworks is an SNPP link. An SNPP link where each subnetwork is in a different layer is known as a transitional SNPP link. It may also be an SNPP link in which the subnetworks are in different sublayers of the same layer. They only occur across boundaries between layers or sublayers where [ITU-T G.800] transitional links can exist.

Figure 7-2 illustrates the relationship between the transport resources described in [ITU-T G.800] and the entities that represent these resources from the perspective of network management as described in [ITU-T M.3100] and [ITU-T G.7711], respectively.



G.7701(16)_F7-2

Figure 7-2 – Relationship of ITU-T G.800 architectural entities with their representation from a management perspective

As may be seen from Figure 7-1 and Figure 7-2 (which provides a simplified version of Figure 5-2 of [ITU-T G.7711]), the logical termination point (LTP) is used to represent transport resources and the FP, forwarding end point (FwEP), and SNP reference points.

The implications on the SNP when it inherits properties from the resource is for further study.

7.4 Virtualization

An abstraction is a representation of an entity in terms of selected characteristics, while hiding or summarizing characteristics irrelevant to the selection criteria. An abstraction may also use identifiers from a name space that is different from the name space of the resources that are being represented. A virtualization is an abstraction and subset whose selection criterion is dedication of resources to a particular client or application. A VN is an abstraction of a designated subset of [ITU-T G.800] layer network resources.

Connection routing in a VN can only use the SNPPs associated with that VN.

An FP can be assigned and allocated to only one VN; in this case the FP is represented by a single SNP (in the VN).

An FP can be assigned (but not allocated) to multiple VNs; in this case the FP is represented by multiple SNPs (one in each VN). Connectivity on a link that is shared between VNs is modelled by creating a (potential) SNP for each of the shared FPs in each VN. When an FP is allocated to a particular SNP in one VN, the SNPs referencing the same FP in other VNs become busy. This is illustrated in Figure 7-3 which shows an example of two VNs, each with two SNPs. The resources support three FPs, FP₁ is assigned to VN2 and is represented by SNP2₁, FP₃ is assigned to VN1 and

is represented by SNP1₁, FP₂ is shared by VN1 and VN2, FP₂ is represented by SNP1₂ in VN1 and SNP2₂ in VN2.

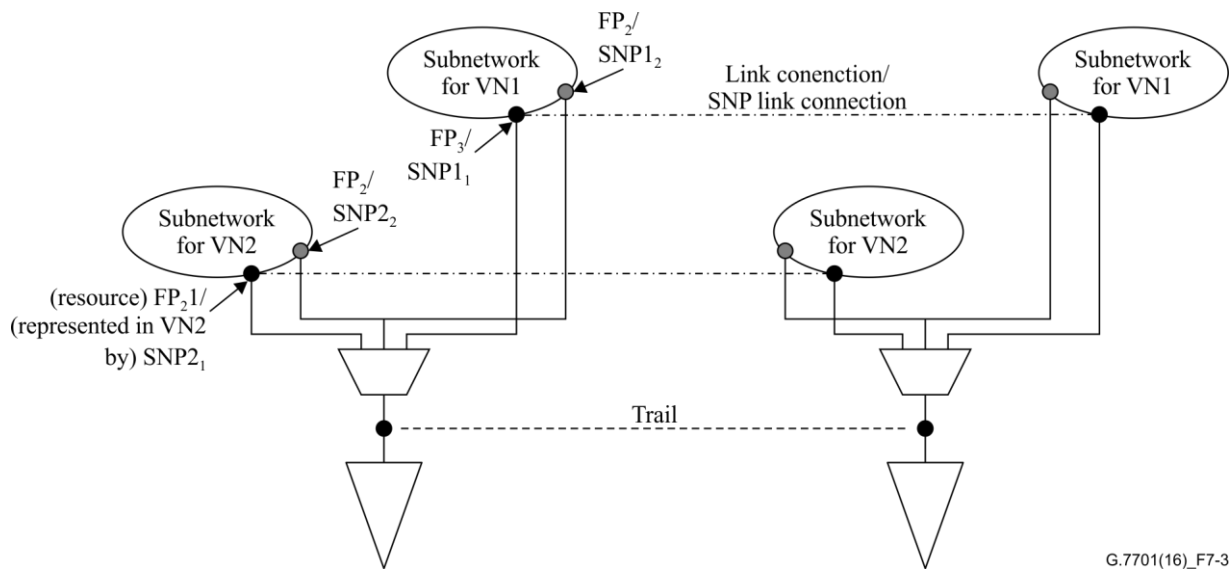


Figure 7-3 – Allocation of link resources between VNs

During the operation of the network if FP₂ is allocated to VN1, SNP1₂ becomes available while SNP2₂ (in VN2) becomes busy.

A resource label is the information required to distinguish a communication within the context of other communications [ITU-T G.800] (an example of a resource label is a multiprotocol label switching (MPLS) label). The SNP references an FP that has a resource label. In the case where the resource label does not carry a resource reservation (e.g., packet switching, see clause 8.3.4), the connection request must include an explicit resource reservation.

8 Control component approach

A control component approach is used to express the architectures for ASON and that for application of SDN to transport in a way that facilitates the construction of reasonable scenarios.

Control components are identified in such a way that the most commonly used distributions of functionality are supported. Interfaces are exposed according to when there is a desire to distribute components among platforms, which is implementation specific. Where this is the case, the primitives exchanged between components are combined into streams carried over an external protocol.

It should be noted that the architecture is not solely comprised of control components and their interactions, but also assumes the presence of key conceptual constructs. A resource database (RDB) is a logical entity that holds information used by components and it is assumed to be present as needed to support components. The RDB is used by various distributions of components and may be used to exchange information between components.

8.1 Notation

In unified modelling language (UML), a component is defined as "a physical and replaceable part of a system that conforms to and provides the realisation of a set of interfaces" [b-UML]. A component in this sense typically represents the physical packaging of otherwise logical elements including classes and interfaces. In the context of this Recommendation, a component is defined as "an element that is a replaceable part of a system that conforms to and provides the realisation of a set of interfaces". A component here represents an abstract entity rather than a piece of implementation code. Thus, in this context, components represent logical functions rather than physical

implementations. With this in mind, UML can be used in describing the component architecture for the ASON and the application of SDN to transport.

This clause considers the component architectural notation based upon some simple building blocks from the vocabulary of the UML.

Component interface: A component interface supports a collection of operations that specify a service of a component, and is specified independently from the components that use or provide that service. Operations specify the information passed in or out together with any applicable constraints. Interface definitions are presented in the form of a table, an example of which is presented in Table 8-1. Each interface has an interface name that identifies the role. Input interfaces represent services provided by the component; the basic input parameters are required for the specific role and basic return parameters are a result of the action on the input parameters. Output interfaces represent services used by the component; the basic output parameters define the information provided, the basic return parameters (if identified) are those required in response to the output parameters. Notification interfaces represent unsolicited output actions by the component, and are represented by an output interface with no return parameters. These three interface types are described separately in interface specifications.

Table 8-1 – Generic component interface descriptions table format

Input interface	Basic input parameters	Basic return parameters
Interface name	Input parameters	Returned parameters

Output interface	Basic output parameters	Basic return parameters
Interface name	Output parameters	Returned parameters

Transaction semantics associated with a particular transaction are assumed to be handled transparently, and there is no need to explicitly mention separate parameters for this purpose in interface description.

Role: A role is the behaviour of an entity when it is participating in a particular context. Roles allow for the possibility that different entities participate at different times, and are denoted by annotating a relationship with the name of an interface.

Component: In this Recommendation, components are used to represent abstract entities, rather than instances of implementation code. They are used to construct scenarios to explain the operation of the architecture. This component is represented as a rectangle with tabs. This is illustrated in Figure 8-1.




Figure 8-1 – Representation of a component

Generically, every component has a set of special interfaces to allow for monitoring of the component operation, dynamically setting policies and affecting internal behaviour, and obtaining resource information. These interfaces are not mandatory, and are provided on specific components only where

necessary. Where appropriate, the use of the monitor interface is described in individual component descriptions. Components are not assumed to be statically distributed.

Component interfaces are defined based on the operations that the component is intended to perform and are, therefore, component specific. When interfaces on components are described, only the different interface types are specified. All components have the property of supporting multiple instances of an interface.

As components are used in an abstract way, this specification is extendable by the techniques of component sub-classing and composition.

8.2 Policy

For the purposes of this policy model, systems represent collections of components, and a system boundary provides a point where policy may be applied. Policy is defined as the set of rules applied to interfaces at the system boundary, and implemented by port controller components. Policy ports are used to simplify the modelling of policies that are applied to multiple ports. System boundaries are nested to allow for correct modelling of shared policies applied to any scope (full system, any set of components, individual components, etc.). The order of the policy application is that which is specified by the nesting.

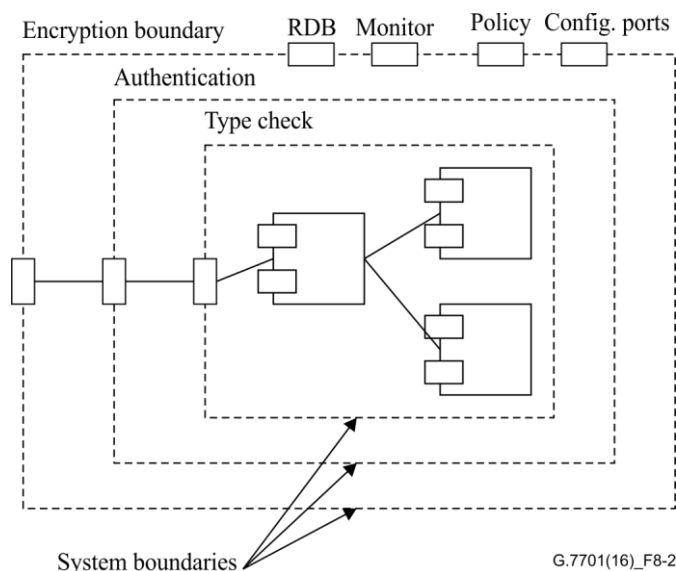


Figure 8-2 – Example of system boundaries related to policy control

In Figure 8-2, the dashed boxes represent system boundaries, while the closed rectangles on the boundary, called ports, represent port controller components.

The monitor, policy, configuration, and RDB ports may be available on every system (and component) without further architectural specification. The monitor port allows management information to pass through the boundary relating to performance degradations, trouble events, failures, etc., for components, subject to policy constraints. The policy port allows for the exchange of policy information relating to components. The configuration port allows for the exchange of configuration, provisioning and administration information relating to components (subject to policy constraints) that may dynamically adjust the internal behaviour of the system. The RDB port allows for the exchange of resource information that is relevant to the component.

Figure 8-2 shows an example of how encryption, authentication and type checking may be implemented as a set of three nested port controllers, where the policy application order follows the nesting order. The components inside the authentication boundary do not specify encryption or authentication requirements, as these are properties of the component environment. Port controllers

are defined for each independent aspect of port policy, and combined policy is achieved by composition of port controllers. This allows the creation of reusable components, which are distinguished by a descriptive prefix. Policy violations are reported via the monitoring port.

The policy port may be seen as a filter of incoming messages, where messages that are rejected have violated the policy. Policies may be dynamically changed via the system policy port, and in this way, dynamic behavioural changes may be described.

It is common to discuss how policy is applied at a reference point, but policy can only be applied to the individual interfaces crossing the reference point. A method of combining several interfaces into a single implementation interface is described later in clause 8.3, which includes a discussion of protocol controllers (PCs).

Other aspects of policy relate to variable behaviour of the components (such as schedules, access rights, etc.) and these aspects are specified and implemented by the components. Component behaviour may also be dynamically changed, and the ability to do this may be controlled by policy. This allows the determination of which aspects of system behaviour are specified where.

Policy, as other aspects of the system, may be distributed. The policy enforcement point (PEP), the point where the policy decisions are enforced, of that model corresponds to the port in this model. The policy decision point (PDP) is the point where policy decisions are made. This can be done within the port, though it may be distributed to a different system. This distribution decision depends on many factors that in turn depend on the actual policy. As an example, performance reasons may force the PDP to be within the port (encryption), while security reasons may force the PDP to be elsewhere (password lookup).

When the PEP and PDP are not collocated, cooperation is required.

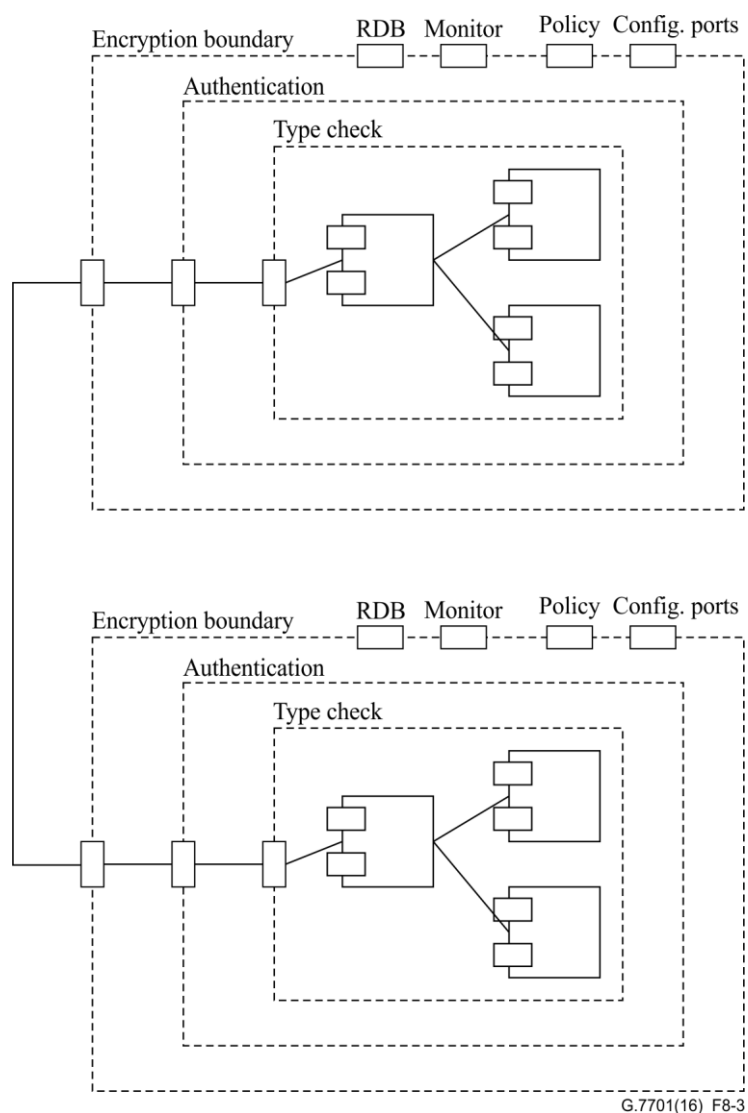


Figure 8-3 – Example of system boundaries related to policy control

Figure 8-3 provides an example of interaction between two different subsystems. They communicate with each other through ports which are connected together. The received messages are decrypted first, then pass authentication type check, and finally processed by the correct control component.

8.3 Common components

The component architecture for the application of control functions to transport networks is described in this clause. A summary of the primary function of each component is first provided, followed by descriptions of component interfaces and detailed operations, respectively.

The connection controller (CC), hierarchical call controller (CallC), and RC are all control components. The VN context includes the name space that control components use to operate on VN resources.

Special components are defined in this Recommendation to allow for implementation flexibility. These components are PCs and port controllers. The detail of the interfaces of these and other components are provided in other technology specific Recommendations.

Protocol controllers are provided to take the primitive interface supplied by one or more architectural components, and multiplex those interfaces into a single instance of a protocol. In this way, a PC absorbs variations among various protocol choices, and the component architecture remains invariant. One or more PCs are responsible for managing the information flows across a reference point.

Port controllers are provided to apply rules to system interfaces. Their purpose is to provide a secure environment for the architectural components to execute in, thereby isolating the architectural components from security considerations. In particular, they isolate the architecture from distribution decisions made involving security issues.

8.3.1 Call controller components

Calls are controlled by means of call controllers, which are able to recurse hierarchically or horizontally.

There are two types of call controller components:

- 1) a calling/called party call controller (CCC): This is associated with an end of a call and may be co-located with end systems or located remotely and acts as a proxy on behalf of end systems. This controller acts in one, or both, of two roles, one to support the calling party and the other to support the called party.
- 2) a network call controller (NCC): This provides three roles, one for support of the calling party, another to support the called party and a third to support calls across domain boundaries.

A calling party call controller interacts with a called party call controller by means of one or more intermediate NCCs.

Network call controllers are instantiated at policy boundaries (e.g., at reference points or aggregated interfaces).

Call controllers that are adjacent (in the context of a call) form a call segment.

The role of this component is:

- processing of incoming call requests;
- generation of outgoing call requests;
- generation of call termination requests;
- processing of call termination requests;
- translation from call source and destination identifiers (which are boundary resource identifiers) to identifiers from the transport resource name space (i.e., SNPs, SNPPs) via directory request;
- call admission control based on validation of call parameters, user rights and access to network resource policy;
- state management of client calls and itself (call state could be held in the RDB);
- application of policies such as quality of service (QoS), protection, access control, and security;
- adaptation management of transport resources via the termination and adaptation performer (TAP) component.

This component has the interfaces provided in Table 8-2 and illustrated in Figure 8-4.

Table 8-2 – Network call controller component interfaces

Input interface	Basic input parameters	Basic return parameters
Call request accept	Boundary resource identifier	Confirmation or rejection of call request
Network call coordination in	Boundary resource identifier	Confirmation or rejection
Call release in	Boundary resource identifier	Confirmation of call release
Client NCC coordination in	Optional client call parameters, optional client context identification, boundary resource identifiers	A pair of SNPs in the client context
Server NCC coordination in	A pair of SNPs	Confirmation or rejection of use
Call modification accept	Call name, parameters to change	Confirmation or rejection of call modification

Output interface	Basic output parameters	Basic return parameters
Call indication	Boundary resource identifier	Confirmation or rejection of call request
Connection request out	Boundary resource identifier	A pair of SNPs
Network call coordination out	Boundary resource identifier	Confirmation or rejection of call request
Directory request	Boundary resource identifier	SNP, SNPP
Policy out	Call parameters	Accept or rejection of call
Call release out	Boundary resource identifier	Confirmation of call release
Client NCC coordination out	A pair of SNPs in the client context	Confirmation or rejection of use
Server NCC coordination out	Optional call parameters, context identification, boundary resource identifiers	A pair of SNPs
Call modification request	Call name, parameters to change	Confirmation or rejection of call modification

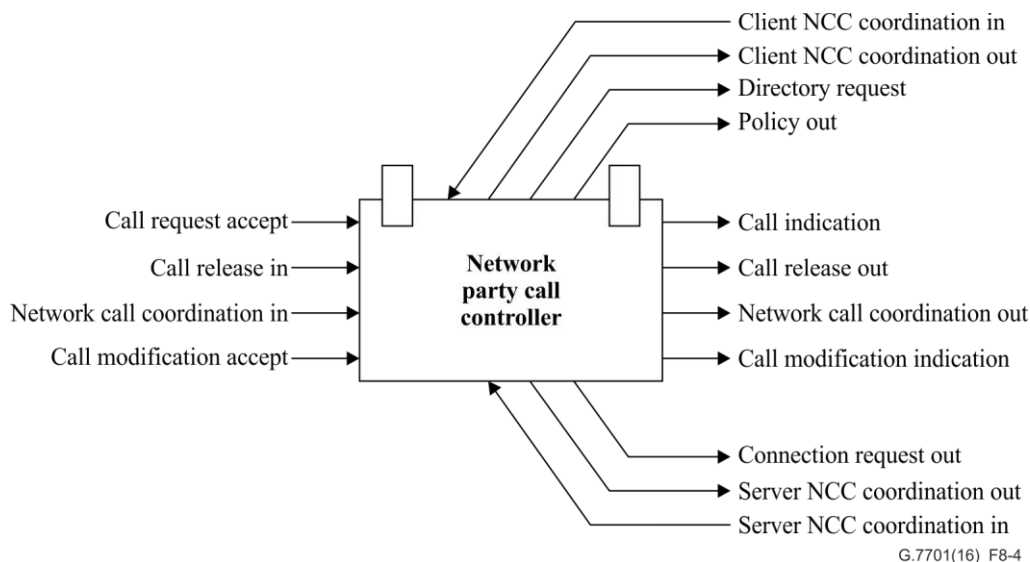


Figure 8-4 – Network call controller component

Call request accept: This interface is used to accept a call source and destination identifier pair. This interface also confirms or rejects the incoming call setup request. It is used in the cooperative model for peer call controllers.

Call indication: Call requests are forwarded to NCCs as calls are progressed. This interface also confirms or rejects the outgoing call setup request. It is used in the cooperative model for peer call controllers.

Connection request out: This interface is used to place a connection setup request to a connection controller as a pair of SNPs.

Directory request: This interface is used to translate a boundary resource identifier (BRI) into an SNP/SNPP identifier.

Network call coordination: This interface is used for network-level call coordination.

Call release in/out: These interfaces are used to place, receive and confirm release requests.

Policy out: This interface provides policy checking.

Client NCC coordination in: This interface is used in NCC hierarchies to accept a request from another NCC for a pair of SNPs. The NCC is provided with source and destination identifiers in its context in order for it to provide a network connection for use by the client context. For NCCs in a client context, SNPs in the client context that are supported by an adaptation to the network connection are returned. For NCCs in the same layer, SNPs are returned that are used by the CC associated with the calling NCC, which is in a hierarchical arrangement. This interface is also used to release or modify the use of the SNP pair. The called NCC returns the result of the action.

Client NCC coordination out: This interface is used in NCC hierarchies to present a pair of SNPs to a client layer that are supported by an adaptation to a network connection, or to an NCC in the same layer that is in a hierarchical arrangement. The client NCC indicates whether or not it accepts this resource. This interface is also used to release or present a modified SNP pair. The called NCC returns the result of the action.

Server NCC coordination out: This interface is used in NCC hierarchies to request a pair of SNPs (input and output) that can be used by the call to transfer CI. For called NCCs in a different layer, it is identical to the return parameters of the connection request out interface except that a network connection in this layer is not assumed to be created. For called NCCs in the same layer, it is identical to the return parameters of the connection request out interface. When the context is in the same layer,

NCCs are in a hierarchical arrangement. This interface is also used to release or request modification of the use of the SNP pair provided by the called NCC. The called NCC returns the result of the action.

Server NCC coordination in: This interface is used in NCC hierarchies to accept a pair of SNPs (input and output) presented from another NCC. It may be accepted or rejected. For the interlayer case, the other NCC is in a server layer. For the same layer case, it is another NCC providing resources in a hierarchical arrangement. This interface is also used by to release or present a modified SNP pair. The NCC returns the result of the action.

Call modification accept: This interface is used to accept a call modification request. This interface also confirms or rejects the incoming call modification request.

Call modification indication: This interface is used to continue a call modification request to another NCC. It also receives confirmation or rejection of the request.

The role of call admission control in the calling party NCC is to check that a valid called user name and service parameters have been provided. The service parameters are checked against a service level specification. If necessary, these parameters may need to be renegotiated with the calling party call controller. The scope of this negotiation is determined by policies derived from the original service level specification, which itself is derived from the service level agreement.

The role of call admission control in the called party NCC, if present, is to check that the called party is entitled to accept the call, based on the calling party and called party service contracts. For example, a caller address may be screened, and the call may be rejected.

The directory request interface of the NCC is used to access a directory function that is used to transform identifiers between or within name spaces. An identifier is supplied as input to the directory function which returns one or more identifiers. How the directory is maintained or configured is outside the scope of this Recommendation and will be described in other Recommendations.

Using the Call request accept and Call indication interface, a call is progressed between CCC and NCCs forming call segments. Using the Network call coordination in and Network coordination out interface, a call is progressed between NCCs forming call segments. A chain of NCCs can be formed that result in a horizontal chain of call segments. For the hierarchical NCC case the Client/Server NCC Coordination in/out interfaces are used. These returned pairs of SNPs represent a connection and allow the calling NCC to gather both call and connection segments.

8.3.2 Connection controller component

The CC is responsible for coordination among the link resource manager (LRM), RC, and both peer and subordinate connection controllers for the purpose of the management and supervision of connection setups, releases and the modification of connection parameters for existing connections. This component services a single subnetwork, and provides the abstract interfaces to other control components given in Table 8-3. The connection controller component is illustrated in Figure 8-5. Connection state could be held in the resource database.

Table 8-3 – Connection controller component interfaces

Input interface	Basic input parameters	Basic return parameters
Connection request in	A pair of local SNP identifiers and optionally a route	A subnetwork connection
Peer coordination in	1) A pair of SNP identifiers; or 2) SNP and SNPP; or 3) SNPP pair; or 4) route	Confirmation signal

Output interface	Basic output parameters	Basic return parameters
Route query	Unresolved route fragment	Route
Link connection request	–	A link connection (an SNP pair)
Connection request out	A pair of local SNP identifiers	A subnetwork connection
Peer coordination out	1) A pair of SNP identifiers; or 2) SNP and SNPP; or 3) SNPP pair; or 4) route	Confirmation signal




Figure 8-5 – Connection controller component

Connection request in: This interface is used to receive a connection request.

Connection request out: This interface is used to initiate a connection request to another connection controller component.

Peer coordination in: This interface is used to receive a connection request that is in progress. Specific or general topology identifiers are included to indicate points that the connection should traverse.

Peer coordination out: This interface is used to continue a connection request that is in progress. It is sent to a peer connection controller.

Link connection request: This interface is used to request a link connection from LRM. It is also used at the bottom of hierarchical recursion and in the peer coordination process.

Connection setup is performed in response to either a connection request, from an enclosing scope connection controller, or from a peer connection controller. In the case of hierarchical routing where the superior (i.e., parent) CC selects the source and destination SNPs, the connection request in/out interface is used. A parent CC may have multiple subordinate CCs in a hierarchy and in that case, the parent CC concatenates the connections together in an end-to-end manner. In all other cases, the peer coordination in/out interfaces are used. Component operation is the same in both cases.

The first unresolved portion of the route is resolved, via the route table query interface, into a set of links to be traversed, and this new set of links adds to the set. The connection controller inspects the new set of links to see which of these links are available for link connection allocation. Link connections are obtained and their links are removed from the link set. Next, corresponding

subnetwork connections are requested from subordinate (i.e., child) connection controllers via the connection request out interface. Any unallocated route components are passed on to the next downstream peer connection controller. The actual sequence of operations depends on many factors, including the amount of routing information available and the access to particular LRMs; however, the operation of the connection controller is invariant. Connection release is an analogous operation to connection setup, except the operations are reversed.

8.3.3 Routing controller component

The routing controller component is an abstract entity that provides the routing function. It can be implemented as a single entity, or as a distributed set of entities that make up a cooperative federation.

The role of the RC is to respond to requests for:

- path (route) information needed to set up connections. This information can range from end-to-end path details to a next hop. The route can be computed by one or more cooperating RCs;
- topology (SNPs and their abstractions) information for control and management purposes.

Information contained in the RC (SNPPs, SNP link connections) enables it to provide routes within the domain of its responsibility. The RC may also receive topology updates from its peers, consisting of identifiers for adjacent nodes/subnetworks and their links, that may also include identifiers of specific peer subnetwork resources. RC information may be contained in the resource database. The resource database may also maintain knowledge of SNP state to enable constraint based routing. Using this view, a possible route can be determined between two or more (sets of) SNPs taking into account some routing constraints. There are varying levels of routing detail that may be provided.

The functional requirements of the RC include:

- topology management for intra-domain and multi-layer transport network;
- topology management for inter-domain and multi-layer transport network;
- path computation of intra-domain, inter-domain or multi-layer for transport network;
- abstraction of topology and resources of transport network.

The component architecture of RC is shown in Figure 8-6.




Figure 8-6 – Component architecture of routing controller

The input and output interfaces for RC are described below.

Table 8-4 – Routing controller component – input interfaces

Input interface	Basic input parameters	Basic return parameters
Route query	Unresolved route element, Route policy	Route
Local topology in	Local topology update	–
Network topology in	Topology update	–

Table 8-5 – Routing controller component – output interfaces

Output interface	Basic output parameters	Basic return parameters
Route query	Unresolved route element, Route policy	Route
Local topology out	Local topology update	–
Network topology out	Network topology update	–

Local topology interface in: The local topology interface in is used to maintain accurate information on the local topology and its state. This is the topology information that is within the domain of responsibility of the RC.

Local topology interface out: The local topology interface out is used to convey information on the local topology and its state.

Network topology interface in: The network topology interface in is used to retain accurate information on the network topology and its state. This is the topology information that is outside the domain of responsibility of the RC. Information from other RCs Local topology interface out and Network topology interface out are received on this interface.

Network topology interface out: The Network topology interface out is used to convey the information received from the Network topology interface in, or a summarized version thereof (further abstracted), and may also provide a summarized version of the local topology

Route query interface: This interface accepts an unresolved route element and returns a route. Route queries may be either from the connection controller component (CC) or may be from another RC.

Depending upon the context, the RC component can receive local topology information from the LRM component or network topology information from an RC at a lower hierarchical level.

In the routing calculation process, RC can cooperate with other RCs to perform multi-domain and multi-layer route calculation.

8.3.4 Link resource manager component

The LRM component is responsible for the management of an SNPP link; this includes the assignment and unassignment of SNP link connections (to a connection), managing resource reservation, the configuration of policing and shaping functions via the TAP (if required), providing topology and status information. The SNPs and the associated resource (e.g., capacity) that are in the SNPP at each end of the link are provided by the TAP (see clause 8.3.6), the LRM also receives resource status information from the TAP. A LRM may request the TAP to modify the allocated capacity or the list of allocated SNP identifiers (i.e., change capacity or the binding state of the SNPs to allocated from potential or from allocated to potential). The LRM functions are shown in Figure 8-7.

Layer network using circuit switching with a fixed bitrate

The TAP supplies the LRM with the set of SNPs in the SNPP. When the TAP allocates an SNP, and associates it with a resource label, the FP and link connection are created; the resource label provides an implicit reservation of the link resource. The LRM may assign any of its SNPs to support a connection without further interaction with the TAP. The LRM can track the utilization of the SNPP link by tracking the SNPs that it has assigned. In general, the same SNP identifier and resource label are used for both directions of a bidirectional connection.

Layer network using circuit switching with a flexible bitrate, (e.g., OTN with ODUflex)

The TAP supplies the LRM with the set of SNPs in the SNPP, the potential and allocated capacity, together with the capacity assignment policy. The FP and link connections are not created (by the TAP) until the LRM assigns an SNP and capacity to a connection. A LRM may request the TAP to modify (i.e., increase or decrease) the capacity used by the connection. The LRM advises the TAP when an SNP is assigned or unassigned, this allows the TAP to create or delete the FP. The LRM is responsible for the assignment of the allocated capacity within the constraints of the capacity reservation policy provided by the TAP. The LRM can track the utilization of the SNPP link by tracking the SNPs and capacity that it has assigned. In general, the same SNP identifier and resource label are used for both directions of a bidirectional connection.

Layer network using packet switching

The SNPP link information held by the LRM must include the admission control policy (e.g., amount of overbooking allowed for the committed information rate (CIR) and peak information rate (PIR)). When the LRM receives a connection create request, or connection modification request, the LRM's connection admission control (CAC) function determines whether the request can be granted or whether it has to be rejected.

The TAP supplies the LRM with the set of SNPs in the SNPP, the potential and allocated capacity (CIR and PIR), together with the capacity assignment policy. The FP and link connections are not created (by the TAP) until the LRM assigns an SNP and capacity (CIR and PIR) to a connection. A LRM may request the TAP to modify (i.e., increase or decrease) the capacity used by the connection. The LRM advises the TAP when an SNP is assigned or unassigned, this allows the TAP to create or delete the FP. The LRM (via the TAP) must also configure the appropriate traffic policing and shaping functions. The LRM is responsible for the assignment of the allocated capacity within the constraints of the capacity reservation policy provided by the TAP. The LRM can track the utilization of the SNPP link by tracking the SNPs and capacity (CIR and PIR) that it has assigned. Different SNP identifiers and resource labels may be used for each direction of a bidirectional connection.

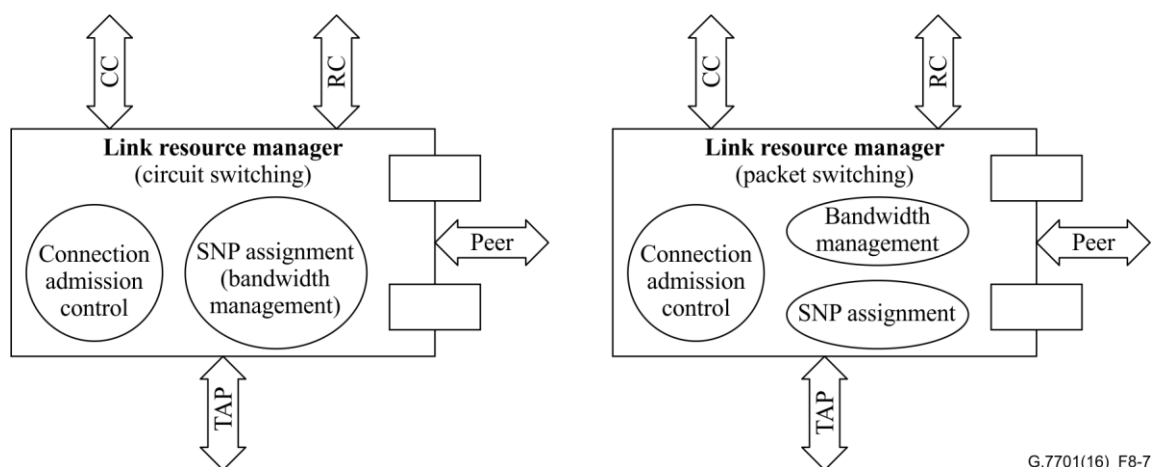


Figure 8-7 – Basic LRM functions for circuit switching and packet switching

The functional requirements of LRM include:

- the output of local topology (nodes and links), resources and abilities information of transport network;
- the status update of topology (nodes and links), resources and abilities information of transport network;
- link configuration, OAM configuration and maintenance;
- assist RC to update the status of topology of transport network;
- assist CC to update the status of connections of transport network.

The component architecture of LRM is shown in Figure 8-8:

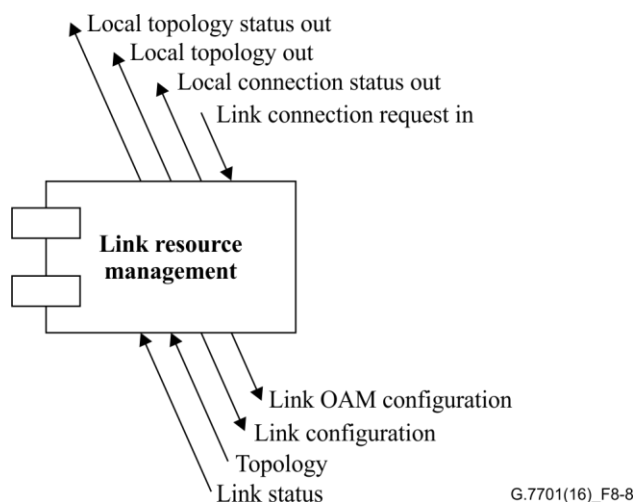


Figure 8-8 – Component architecture of LRM

Table 8-6 – Link resource management component – input interfaces

Input interface	Basic input parameters	Basic return parameters
Link connection request	A link connection (an SNP pair)	Confirmation or rejection
Topology in	Local topology update	–
Local link status in	Link connection status	

Table 8-7 – Link resource management component – output interfaces

Output interface	Basic output parameters	Basic return parameters
Local topology out	Local topology update	–
Local topology status out	Network topology update	–
Link OAM configuration	OAM parameters	Confirmation or rejection of OAM parameters
Link configuration	Link information	Confirmation or rejection of link information
Local connection status out	Local connection status <ul style="list-style-type: none"> – operational state – administrative state – performance 	

8.3.5 Discovery agent component

The discovery agent (DA) component deals with transport resources that have not yet been assigned to a control function. The federation of discovery agents operates solely in the transport resource name space, and provides for separation between that space and the control name spaces. The federation has knowledge of FPs in the network, while a local DA has knowledge of only those FPs assigned to it.

The transport resource name space provides the native identifiers used by transport resources; this Recommendation assumes that a transport resource name space exists for the ITU-T G.800 FPs in the layer network. SNP name spaces are mapped to those FPs and the mappings are maintained by TAP components. The TAP also maintains the relationship between the client forwarding point (cFP) and server forwarding point (sFP).

DAs discover the trail FP to FP relationship from which the link relationships from client FP to client FP are inferred [ITU-T G.7714.1]. Generally, multiple client FPs (cFPs) are associated with a trail FP in the server layer (sFP). Client FPs result from adaptation and if the adaptation is flexible, FPs in different layers may result. LRMs hold the corresponding SNP to SNP relationship for links in the SNP name space.

FP identifiers are only used in the context of DA and TAP components and all other control components use SNP identifiers.

The DA uses the coordination in/out interface to communicate with adjacent DAs to obtain the sFP(s)-sFP(s) relationship between the local sFP and the remote sFP. The local cFP(s) and sFP and the remote cFP(s)-sFP binding relationship is provided by the local and remote TAP. With this information, the cFP-cFP relationship can be determined.

A DA can be located on any suitable platform. That is, it can be on a network element (NE), the lowest level controller, or any platform that is independent of NE and controller.

Figure 8-9 depicts actions of the DA used to establish the cFP-cFP relation.

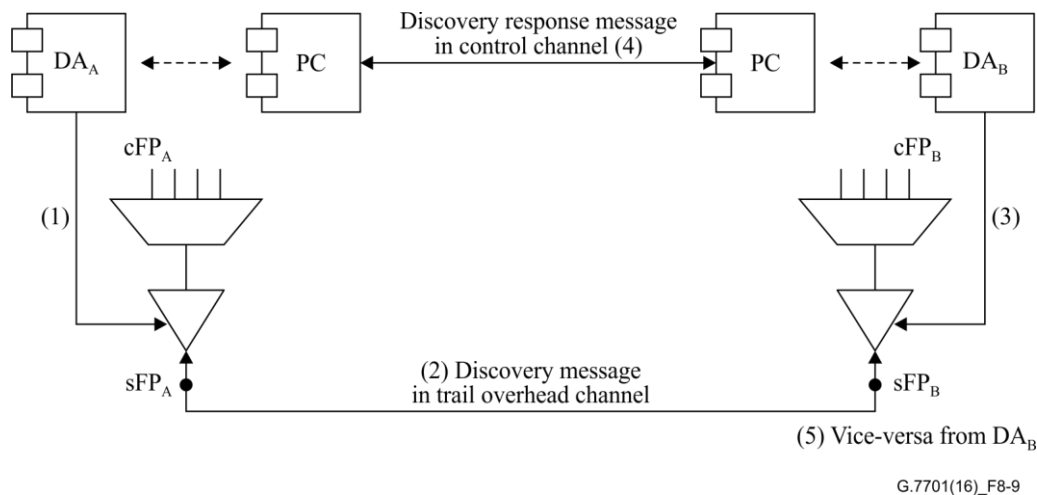


Figure 8-9 – Discovery agent actions

The steps to establishing FP to FP relationships are:


- 1) DA_A triggers sFP_A to inject a discovery message (DM) that contains the sFP_A identifier and the DA_A identifier.
- 2) The discovery message is carried by the trail from sFP_A to sFP_B.
- 3) DA_B receives the discovery message from sFP_B, which contains sFP_A's identifier and the DA_A identifier.

- 4) DA_B sends a discovery response message to DA_A that provides the pair of server FPs (sFP_A and sFP_B) and the DA_B identifier.
- 5) The same discovery process (steps 1 through 4) is performed by DA_B in the reverse direction.
- 6) When the discovery processes in the two directions have completed, each DA locally compares the discovered sFP relationships in the forward and reverse directions to determine whether the relationships in the two directions $A \rightarrow B$ and $B \rightarrow A$ are consistent and the sFP_A – sFP_B relationship has been successfully discovered.

Table 8-8 – Discovery agent component interface

Input interface	Basic input parameters	Basic return parameters
Coordination in	Receives DA response message (responder DA identifier and pair of sFP identifiers)	
DM received	Initiator sFP identifier and initiator DA identifier	

Output interface	Basic output parameters	Basic return parameters
Coordination out	Sends DA response message (responder DA identifier and pair of sFP identifiers)	
DM sent	Initiator sFP identifier and initiator DA identifier	



G.7701(16)_F8-10

Figure 8-10 – Discovery agent component

Coordination in: The coordination in interface receives the DA discovery response message that includes the responder DA identifier and the pair of sFP identifiers.

DM received: This interface receives the in-band discovery message that was sent by the initiator sFP that includes the initiator sFP identifier and the initiator DA identifier.

Coordination out: The coordination out interface sends a discovery response message including the responder DA identifier and pair of sFP identifiers.

DM sent: The initiator DA uses this interface to trigger the initiator sFP to inject an in-band discovery message that contains the initiator DA identifier and the initiator sFP identifier.

Specifics of the discovery message are dependent on specific technology but there are no differences in the discovery procedure for circuit-switched networks, connection oriented packet switched networks and connectionless oriented packet switched networks.

8.3.6 TAP component

8.3.6.1 Termination and adaptation performers

The TAP is collocated with the resources (i.e., the adaptation and termination functions). It provides the LRM with a view of the resources that support a link end; it abstracts the hardware and technology with specific details of the adaptation and termination functions. When utilized in an SDN control context, the TAP component resides at the bottom of SDN controller recursion.

8.3.6.2 TAP resource model

Only those transport resources that are allocated to a control function are made visible to the TAP. When a resource is permanently withdrawn (unallocated) from a control function, the SNP that references that resource should be deleted. The LRM is responsible for the assignment of the allocated capacity within the constraints of the capacity reservation policy provided by the TAP. An LRM may request the TAP to modify the allocated capacity or the list of allocated SNP identifiers (i.e., change capacity or the binding state of the SNPs to allocated from potential or from allocated to potential).

The relationship between the TAP, LRM and the network resources is shown in Figure 8-11.

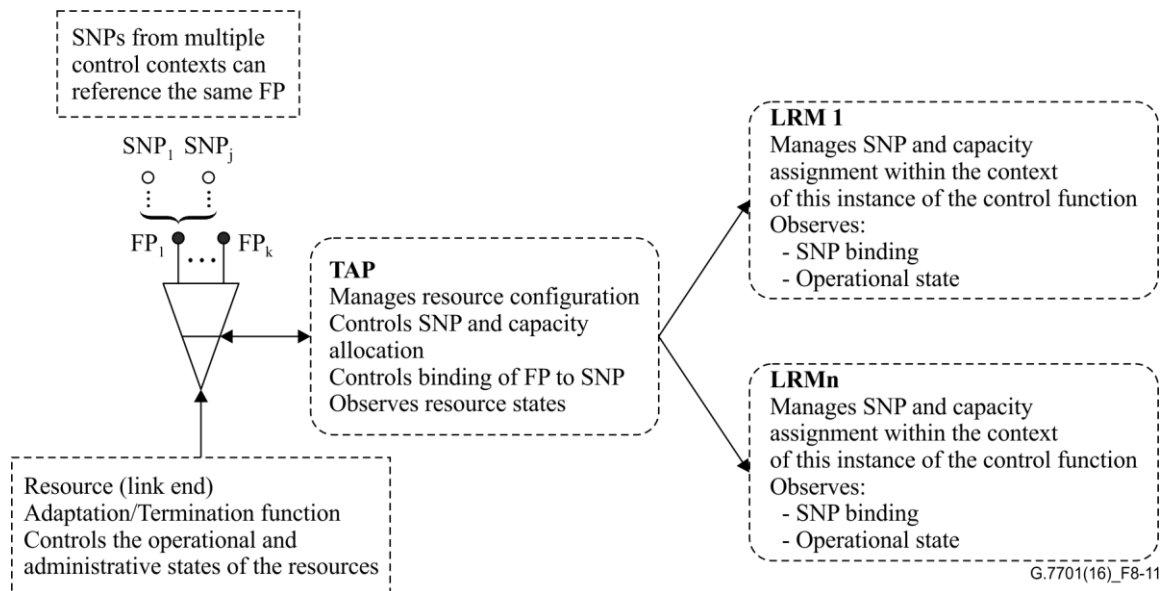


Figure 8-11 – Relationship between TAP, LRM, and transport resources

NOTE – The multiple control contexts may exist because of different VNs or different levels in a controller hierarchy, or both.

The TAP operates at two different times and provides two different functions.

When a resource is assigned to a control function, the TAP is configured with a list of the resource identifiers, the capacity of the link resource, together with the capacity reservation policy. The link resources may be shared between multiple control functions (e.g., different layer networks or different VNs, see clause 7.4). For each LRM that is within the scope of the TAP (i.e., references resources controlled by the TAP), the TAP is configured with the permitted bindings between the resource labels and SNPs. The TAP controls the allocation of SNPs and capacity to each LRM. In a circuit-switched network with a fixed bitrate, only the resource labels are configured since they carry an implicit resource capacity and reservation policy. In the case of a time division multiplexing (TDM) network with flexible bitrate (e.g., OTN with ODUflex) the TAP provides the potential and allocated capacity to each LRM, together with the capacity assignment policy. In the case of packet-switched networks, the TAP provides the potential and allocated capacity (CIR and PIR) to each LRM, together

with the capacity assignment policy. The LRM can only assign SNP link connections (to a connection) if the SNPs and resource capacity has been allocated by the TAP.

The TAP makes resources visible to an LRM by associating a resource label to an SNP identifier. The existence of the SNP identifier is independent of the configuration of the resources.

In the case of a circuit-switched network with fixed capacity, when the TAP allocates the resource to an LRM (i.e., sets the SNP state to allocated), it also configures those resources and creates the FP and the link connection is created. This configuration action is independent of the assignment of those resources to a connection.

In the case of a circuit-switched network with flexible capacity, when the TAP allocates capacity and SNPs to a LRM (i.e., sets the SNP state to allocated), it only performs the configuration required to allow those resources to be activated. When the LRM assigns an SNP to a connection, the TAP creates the FP (i.e., it activates the binding between the SNP and the FP), it also configures the resource to support the requested capacity.

In the case of a packet-switched network, when the TAP allocates capacity and SNPs to an LRM (i.e., sets the SNP state to allocated), it only performs configuration required to allow those resources to be activated. When the LRM assigns an SNP to a connection, the TAP creates the FP (i.e., it activates the binding between the SNP and the resource label), it also configures the traffic conditioning and policing functions, if required.

Figure 8-12 below depicts the relationships between the potential and allocated resource identifiers, or resource labels, and SNPs. Moreover, it shows that the assigned SNPs are associated with labels from the allocated label range that are configured labels, i.e., a link connection exists for those SNPs. For a given label in the transport resource name space, the pair of SNP identifiers allocated for it are the same on both ends of the link. That is, $S_{nA}=S_{nZ}$, for $n=1..m$ in Figure 8-12.

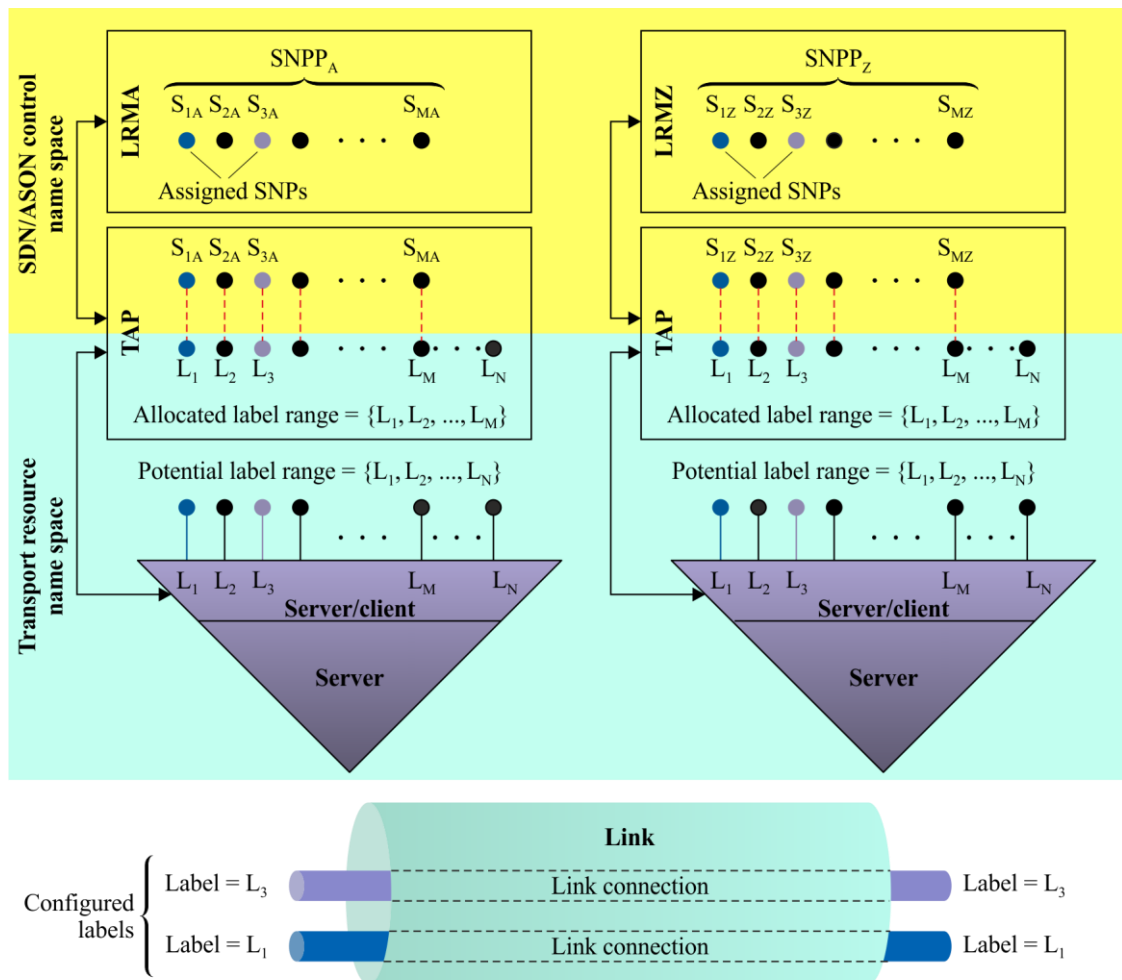


Figure 8-12 – Control components and link resource model

The various types of resource labels are:

– **Potential (resource) label range**

The "potential label range" is the full label range of resource labels in the transport resource name space that an adaptation function supports. In packet switching layers, this range can be much larger than the allocated label range. Example: the 20-bit MPLS label provides $220 = 1048576$ possible label values including reserved labels (label values 0..15) for specific purposes.

– **Configured (resource) label**

A "configured label" is a transport resource label that has been configured in support of a connection. If a label is configured, a forwarding table entry exists on the receiving end of the link such that packets can be forwarded to an outgoing link if a packet is received with a label value that is equal to the configured label. If a label is configured, a packet flow can be distinguished from other flows and can be forwarded based on the label value. This is equivalent to the existence of a link connection. This means that a link connection is created whenever a label has been configured consistently on either end of a link. The deletion of the configuration entry also deletes the link connection.

– **Allocated (resource) label range**

The "allocated label range" is the set of labels that can be used by the adaptation function of a particular link to carry user traffic. It is a subset of the potential label range. The allocated label range must not include reserved label values. When a system uses a per platform (system) label space, each interface is typically configured with an (allocated) label range that does not overlap with the label

ranges of the other interfaces, and a specific label value is selected from this label range in response to, e.g., a connection request.

The allocated labels are entities that can be referenced in the transport resource name space. Each allocated label is associated with one or multiple SNP identifiers that exist in the control component name space (1:n relationship). In the simplest case, there is exactly one SNP identifier per allocated label (1:1 relationship between allocated label and SNP identifier). TAP holds the binding information between SNPs and an allocated label.

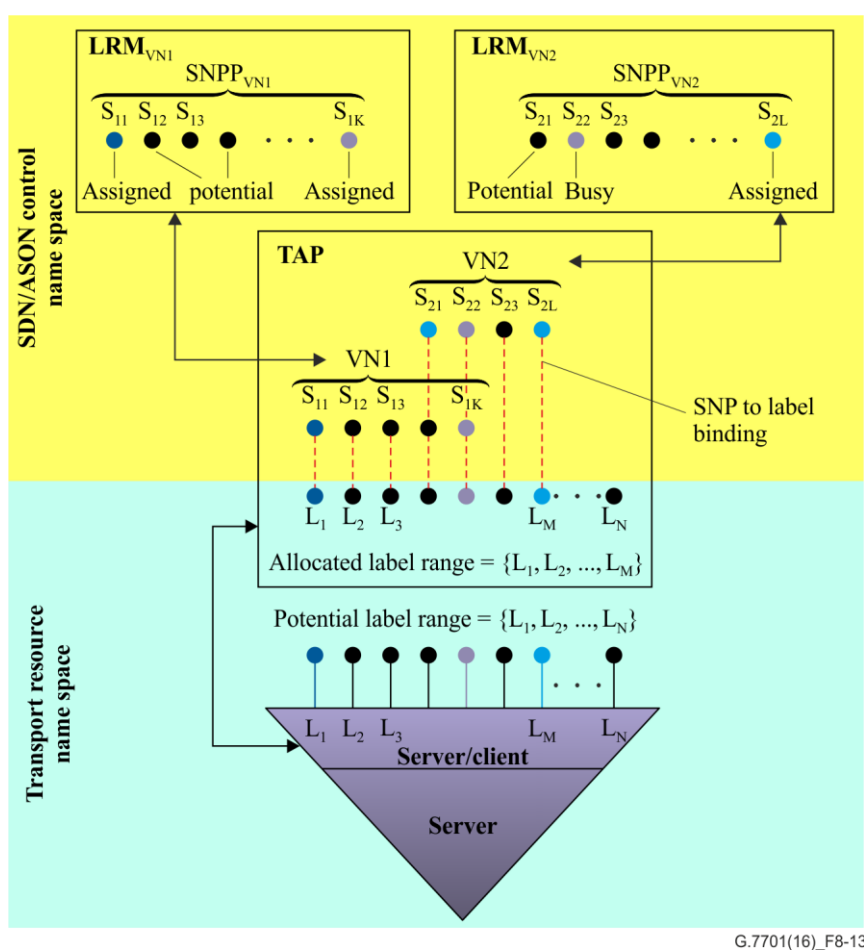
– Potential SNPs

Potential SNPs are those SNPs that are associated with a label. In general, multiple SNPs can be associated with a single label.

– Assigned SNPs

Assigned SNPs are those SNPs out of the set of potential SNPs that have been assigned to a particular connection. This means that the associated label is a configured label.

Figure 8-13 depicts how VNs can be modelled. In the provided example, the allocated label range is subdivided into three subsets, a subset that can exclusively be used by VN1, a subset that can exclusively be used by VN2, and a subset that is shared between VN1 and VN2. Each label of the subsets for exclusive use has a single SNP associated, whereas each label of the shared subset is associated with two SNPs, one SNP within the scope of LRM_{VN1} and another one SNP within the scope of LRM_{VN2} . When an SNP is assigned, e.g., by LRM_{VN1} that corresponds to a label from the shared label subset, the SNP in LRM_{VN2} becomes busy.



G.7701(16)_F8-13

Figure 8-13 – Control component organizational model for VNs

8.3.6.3 TAP states

The TAP holds the SNP binding states and the capacity allocation to each LRM, and provides a specific (coordinated) view to each LRM. As described in Table 8-9, the SNP binding states that the TAP provides to the LRM are constrained by the administrative state of the resources.

The transport resources are aware, from the [ITU-T X.731] usage state (idle, busy), if the resources have been allocated to the TAP. The resources have no visibility of any allocation that the TAP makes to LRMs. Therefore, the resources should use the shutting down state to withdraw resources from the TAP.

The TAP uses the SNP binding states to allocate resources to LRMs. The TAP has no visibility of the assignment of those resources to connections. Therefore, the TAP should use the SNP binding state of shutting down to remove resources from the LRM.

Table 8-9 – SNP binding states

State	Description
Busy	Permitted binding, the resource label and capacity being referenced by the SNP is currently allocated to another control function or the management function.
Potential	Permitted binding, currently the resource label and capacity being referenced by the SNP is not allocated to any control function or the management function.
Allocated	Permitted binding and the resource label and capacity being referenced by the SNP has been configured for and allocated to this LRM.
Shutting down	TAP notification that the resource label and capacity being referenced by the SNP must be returned within an explicit timeframe e.g., <ul style="list-style-type: none">– immediately (interrupt the current call)– quickly (re-route call before dropping)– next maintenance window– when call is dropped

When an SNP identifier is in the allocated state, the TAP must correctly configure the resources (e.g., variable adaptation) and set the state of any other SNPs referencing the same resource to busy.

When SNP identifiers are bound to their corresponding FP, the TAP is responsible for holding the SNP-FP binding. A local TAP cooperates with a remote TAP via the LRM to coordinate any variable adaptation or other coordination required when forming the FP link connections.

If an LRM wishes to use capacity or an SNP with a binding state of "potential" to satisfy a connection request then during connection setup, a pair of TAPs cooperate via the LRM to coordinate any adaptation setup, or link resource allocation, required by the link connection.

When the TAP modifies the resource capacity that is allocated to an LRM, it also makes a corresponding adjustment to the potential resource capacity.

The TAP provides SNP state information to the LRM and accepts resource state status from the adaptation and termination functions to ensure that the management function indications are consistent. Management function consistency includes ensuring that the alarm state of the link connection is consistent, so that spurious alarms are neither generated nor reported.

There are three [ITU-T X.731] states for transport resources:

- 1) operational: This state reflects the combined status of the trail supporting the link and adaptation function. It is controlled by the underlying resources and is observed by TAP.
- 2) administrative: This state reflects the permission to use the resource which is managed by a management interface to the TAP.

- 3) usage: This state reflects whether the resource is actively in use. As TAP allocates and unallocates resources to the control function, it adjusts the usage state accordingly.

Permitted combinations of the resource states and the SNP binding state for each SNP are described in Table 8-10 below:

Table 8-10 – Resource and SNP binding states

ITU-T X.731 resource states			SNP binding states	
Operational	Administrative	Usage	LRM x	All other LRMs (Note 2)
Enabled, disabled (Note 1)	Unlocked	Busy	Potential	Potential
Enabled, disabled (Note 1)	Unlocked	Busy	Allocated (Note 3)	Busy
Enabled, disabled (Note 1)	Shutting down (Note 4)	Busy	Shutting down	Busy
Enabled, disabled	Locked (Note 5)	Idle	Busy (Note 6)	Busy (Note 6)
<p>NOTE 1 – When an LRM observes that the operational state of a link is disabled, it may notify the RC component, it may also notify the connection controllers for the connections that are impacted. The call controller manages the recovery of any connections that are using a failed link.</p> <p>NOTE 2 – If an LRM does not contain an SNP that references the same resource, then the binding state is not present.</p> <p>NOTE 3 – The LRM assigns allocated SNPs and resource capacity to a connection. These assignments are not visible to the TAP.</p> <p>NOTE 4 – If the resource administrative state is changed from unlocked to shutting down, then the TAP must change the binding state of any allocated SNPs that are referencing that resource to shutting down.</p> <p>NOTE 5 – If the resource administrative state is set to locked, then the TAP must set the SNP binding state to busy.</p> <p>NOTE 6 – This combination occurs when the resource is allocated to the management function or when the resource is being withdrawn from the control function. The management function will operate directly on the transport function resources. Changes to the [ITU-T X.731] states will not be visible to the TAP during this time.</p>				

8.3.6.4 Adding/removing resources from a control function

The resource administrative and usage states may be used to control the addition or withdrawal of a resource from the control function. This is illustrated in the administrative state transition diagram in Figure 8-14.

This is visible to TAP and is the [ITU-T X.731] resource state.




Figure 8-14 – Resource administrative state

Table 8-11 – Transition descriptions

Transition	Description	Interface
1	Resource is made visible to the TAP	Usage update
2	TAP is permitted to use the resource The usage state is set to busy	Resource state
3	The resource is being withdrawn	Resource state
4	TAP has set the binding state of all SNPs that reference the resource to busy	Usage update
5	Resource is withdrawn from the control function: If the withdrawal is permanent, then the TAP is instructed to delete all SNPs that referenced the resource	Resource state

8.3.6.5 SNPx binding state transitions per LRM

Figure 8-15 shows the SNP binding state held by an LRM. This is the view that the TAP provides to each LRM based on the [ITU-T X.731] state of the resources. Operations on the TAP affect each LRMs' SNP binding state view.

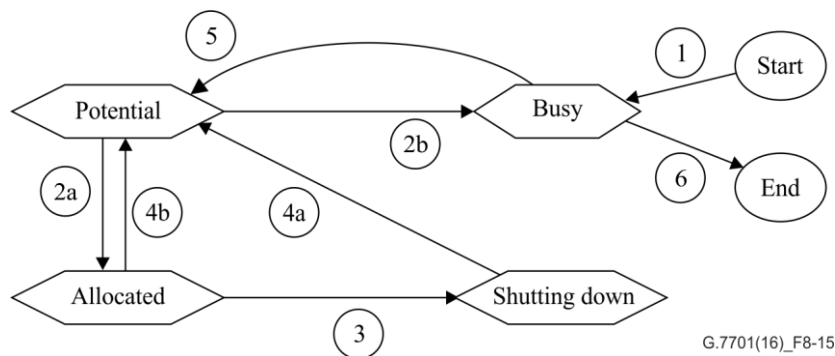


Figure 8-15 – SNP binding state in LRM

Table 8-12 – Transition descriptions

Transition	Description	Interface
1	TAP adds SNP in the scope of the LRM	Add SNP
2a	TAP allocates resource to an LRM	SNP binding state; SNP operational state
2b	TAP sets the SNP binding state to busy when: a) the TAP has allocated the resource to another LRM; or b) the administrative state of the resource has been set to shutting down	SNP binding state
3	TAP requests return of a resource	SNP binding state
4a	LRM is no longer using the resource TAP modifies states to potential	Release SNP
4b	LRM is no longer using the resource TAP modifies states to potential	Release SNP
5	TAP moves resource to potential since it is: a) no longer allocated; or; b) the administrative state has been set to unlocked	SNP binding state
6	SNP is removed from the scope of the LRM	Withdraw SNP

8.3.6.6 TAP component interfaces**Table 8-13 – TAP component interface**

Input interface	Basic input parameters	Basic return parameters
Operational state	Enabled, disabled	Confirm
Administrative state	Locked, unlocked, shutting down	Confirm for locked, unlocked User quit for shutting down
SNP id assigned/unassigned (packet switched only)	SNP id (from LRM) CIR and PIR	
Capacity change request	List of SNP ids CIR and PIR (packet switched only)	Link configuration

Output interface	Basic output parameters	Basic return parameters
Control	Hardware specific	Hardware specific
(Link) Configuration	List of SNP ids CIR and PIR, capacity assignment policy (packet switched only)	confirm
Capacity change (packet switched only)	CIR PIR	Confirm
SNPx binding state	Busy, potential, allocated, shutting down	Resource released (in response to the shutting down state)
SNPx operational state	Enabled, disabled	Confirm
Add SNP	List of SNP identifiers	Confirm
Withdraw SNP	List of SNP identifiers	Confirm
Usage update	New user, user quit	Usage state (idle, busy)

Operational state: This interface accepts resource state information from adaptation and termination functions.

Administrative state: This interface accepts administrative state.

SNP id assigned/unassigned: This interface receives notification of SNP binding actions from LRM.

Capacity change request: This interface receives requests from LRM to change the capacity of packet resources associated with its assigned SNPs.

Control: This hardware specific interface allows the TAP to communicate with the resources that it controls.

Configuration: This interface allows the TAP to provide the link end configuration information to an LRM.

Capacity change: This interface is used by the TAP to advise the LRM if the capacity of the link has been modified. This interface is only used for packet switching.

SNPx binding state: SNP binding state is sent to an LRM.

SNPx operational state: SNP operational state is sent to an LRM.

Add SNP: This interface is used to inform an LRM of a new SNP.

Withdraw SNP: This interface is used to inform an LRM of the removal of an SNP.

Usage update: This interface provides resource state usage information to adaptation and termination functions.

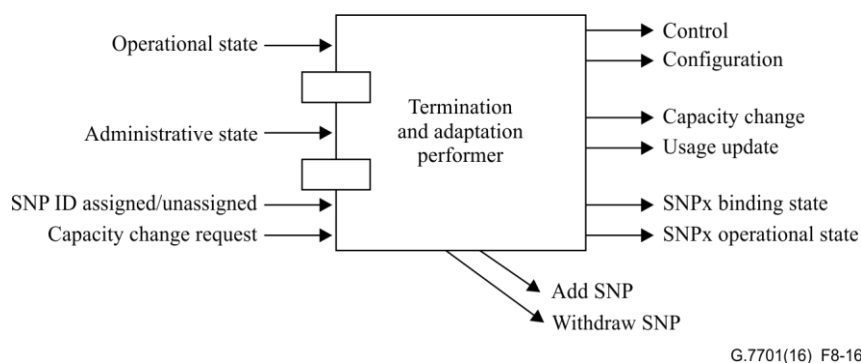


Figure 8-16 – Termination and adaptation performer component

8.3.7 Directory service component

The directory service component is responsible for identifier resolution and coordination among peer directory service components. The role of this component is to provide mappings between name spaces for other components; for example, mapping between VN name spaces.

9 Common control communications

Various applications (management, SDN, ASON, overhead communications network (OCN), etc.) require a communications network to transport information between a variety of components, as illustrated in Figure 9-1. [ITU-T G.7712] specifies DCN functions that can be used to support one or more application's communication network (e.g., communications among management components, communications among control components, and other operations communications).




Figure 9-1 – Example applications supported by a DCN

10 Common management aspects of common control components

This clause introduces constructs relevant to the management of control components.

10.1 Control component relationships

Figure 10-1 illustrates the relevant control components as defined within previous clauses of this Recommendation.

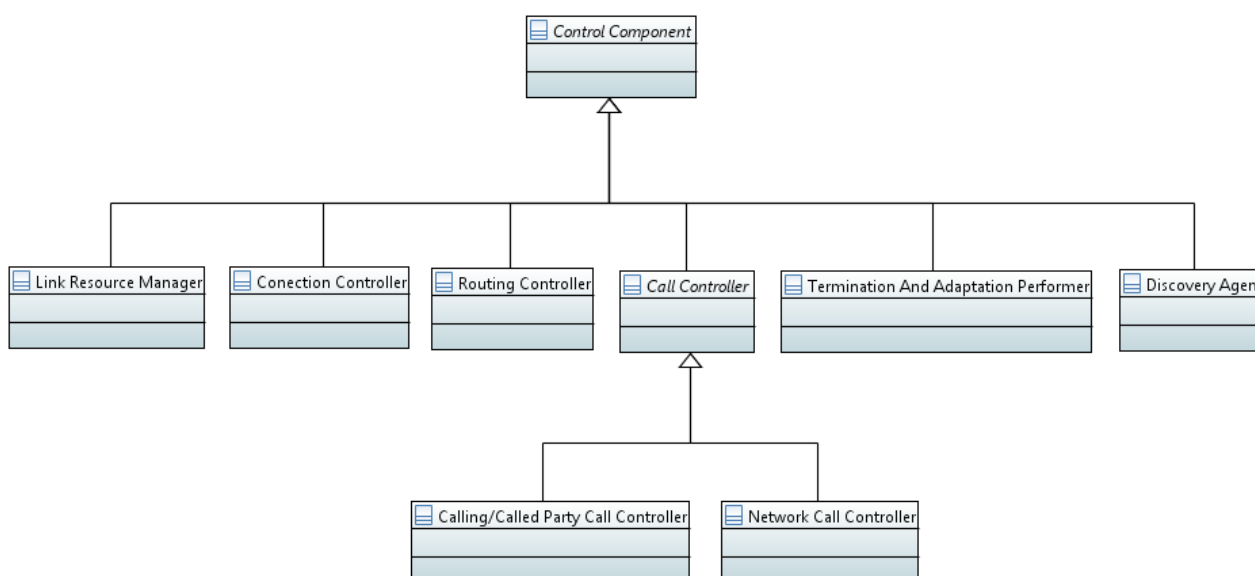


Figure 10-1 – Control component relationships

The following management functions apply to the control components shown in Figure 10-1. Accounting management and security management requirements are for further study.

- TAPs require fault management, configuration management, and performance management.
- DAs require fault management, configuration management, and performance management.
- LRMs require fault management, configuration management, and performance management.
- NCCs require performance management including call statistics, e.g., number of call completed, number rejected, etc. NCCs also require fault management and configuration management.
- RCs require fault management, configuration management, and performance management.
- CCs require fault management, configuration management, and performance management.

10.2 Control function management requirements

The fundamental requirements for control function management are for further study.

11 Identifiers

Distinct and independent sets of name spaces exist, from which identifiers are drawn, for:

- resources in the transport network;
- control view of transport resources;
- control components;
- control artefacts;
- reference points;
- control communications network.

11.1 Resources in the transport network

The architecture of transport resources is described in [ITU-T G.800] and summarized in clause 7.1. The identifiers used by transport resources to allow it to deliver communications from a source to a defined set of destinations are described in [ITU-T G.800] (clause 7, Annex A.4 and Appendix I). [ITU-T G.800] describes topological components that are relationships between reference points, for example, the subnetwork and link. Transport processing functions in the architecture manipulate information at points, and include functions such as adaptation and termination. Transport entities provide the means to transfer information and include connections.

A topological component may have zero or more identifiers associated with it. Similarly, transport processing functions and transport entities may have zero or more identifiers associated with it. A common use of multiple identifiers for the same component is when multiple applications refer to the same component. This Recommendation assumes that there is a singular [ITU-T G.800] FP name space for a set of resources in a layer network.

11.2 Control view of transport resources

11.2.1 Name spaces for routing and connection control

Components that perform routing control operate on a particular set of transport resources within the scope of their designated routing domain/routing area. As defined in clause 3.2, a routing area is defined by a set of subnetworks, the SNPP links that interconnect them, and the SNPPs representing the ends of the SNPP links exiting that routing area. Thus, components collectively use three separate name spaces in order to reference routing area and SNP/SNPP abstractions. These are the routing area name space, subnetwork name space, and link context name space.

The link context name space specifies within the SNPP where the SNP is. At some (low) level in a hierarchy of controllers the SNP identifier must be mapped to a resource label and the TAP component can maintain this mapping. The TAP makes resources visible to an LRM by binding a transport resource label to an SNP identifier.

An SNPP identifier is a concatenation of:

- one or more nested routing area identifiers;
- an optional subnetwork identifier within the lowest routing area level; this can only exist if the containing RA identifiers are present;
- one or more nested resource context identifiers.

This Recommendation does not specify formats or values for identifiers in instances of the aforementioned name spaces. It is possible for implementations to assign the same value to identifiers in different name spaces (e.g., SNPP identifier and BRI) in order to simplify configuration.

11.2.2 Name space recursion

An SNPP identifier can recurse with routing areas as the routing area name space is typically hierarchical. As recursion is a property of [ITU-T G.800] subnetworks it would be natural, though not mandatory, for the naming associated with subnetworks to also be hierarchical.

[ITU-T G.800] also describes recursion between layers using adaptation and termination functions. While name spaces are independent between layers, if an operator owns and manages a client layer supported by a server layer, it is a common operational requirement to understand how a server layer trail can support a client layer connection and in this case a common name space may be used.

Some relationship between recursive name spaces thus needs to be maintained and one mechanism for managing such name space interactions is through the use of the directory service component.

11.3 Control components

Control components also require separate name spaces, as they may be instantiated differently from each other for a given control function instance. For example, one can have a centralized NCC with distributed CCs. Thus, separate identifiers are needed for RCs, NCCs, and CCs. Additionally, the PCs that are used for protocol specific communication also require a separate name space.

11.4 Control artefacts

Control functions create and use artefacts including, for example, connections, routes, calls, and directories. These artefacts have identifiers associated with them.

11.5 Reference points

At the boundary of administrative domains, the BRI is used to identify both the transport resources and the interface between control components that are exposed at the domain boundary. The ITU-T G.8080 user-network interface (UNI) is an example of a reference point.

11.6 Control communications network

To enable control components to communicate with each other, a control communications network (CCN) is used as specified in [ITU-T G.7712]. CCN addresses identify the points of attachment for the PCs that instantiate control communication functions (generating and processing messages in protocol specific formats). Each PC has a distinct CCN point of attachment. For example, the routing PC CCN address refers to the point where the routing PC attaches to the CCN. Thus, the routing PC CCN address is based on the topology of the CCN carrying routing messages, rather than the topology of the transport network or control components.

Annex A

Configuration of OTN digital and media layers

(This annex forms an integral part of this Recommendation.)

[ITU-T G.872] describes the OTN digital layer networks and the media layer. The abstractions used to represent the digital layer networks and the media are similar. However, configuration of the media and the association of media channels to optical signals requires some special consideration.

The configuration of a digital layer network may be simplified since the hypothetical reference model used to allocate impairments has sufficient scope (and implementations have sufficient operating margins) that impairments do not need to be considered. Further, layer independence allows a client (digital) layer network to be configured in a way that is largely independent of the server (digital) layer network.

This simplification cannot be applied to all cases of configuring the media and optical tributary signal (OTSi). As a result, further requirements are imposed on a media control function¹ that is intended to configure the media layer and the signals that cross that media layer.

- The media can be configured in advance of the assignment of the media channel to one or more OTSi
 - The configuration of a media channel is independent of the presence (or absence) of any optical signal.
- The width of the media channels may be different from the width required to carry a single OTSi. (e.g., a "wide" media channel may be used to provide what is commonly called an "express" channel in a ROADM)
- A single client digital stream (optical transport unit (OTU)) may be carried by more than one OTSi
 - Hence the configuration of the media to carry an OTU may require the configuration of one or more network media channels. Further, some of the optical tributary signal assembly (OTSiA) OAM functions are supported by non-associated overhead (OTSiG-O) and this must be configured when an OTSi is assigned to a network media channel.
- Determining the compatibility between an OTSi and a network media channel is a complex process. There is currently no standardized method defined for confirming compatibility in the case where the media is divided across dense wavelength division multiplexing (DWDM) line segments from more than one vendor as described in [ITU-T G.680] scenario 2. Three single domain cases must be considered:
 - 1) the black link approach as described in [ITU-T G.698.1] and [ITU-T G.698.2]. In this case, compatibility can be checked by considering the application codes defined in [ITU-T G.698.1] and [ITU-T G.698.2].
 - The optical domain is normally configured by vendor specific media control, which would not need to manage the optical signal terminations.
 - 2) pre-computed optical paths: In this case paths through the media network are pre-computed and only those that are compatible with the candidate OTSi are provided in response to request for a network media channel. In general, the determination of compatibility is vendor specific.

¹ The media control function may be implemented by, for example, an EMS/OSS, an SDN controller or an ASON/GMPLS control plane.

- 3) compatibility check at the time of request: In general, the determination of compatibility is vendor specific. In this case paths through the media are computed, and after path computation is complete these paths are checked for compatibility with the OTSi. Compatibility checking is a complex process that must take into account the interaction of the "new signal" with all existing or planned² signals that share the same fibre and the impact of these existing (or planned) signals on the candidate "new signal". The actual transfer characteristics of the media elements in the path (e.g., filters, amplifiers) also needs to be considered.

A media³ control function must, in addition to the capabilities required to manage digital layer networks, provide the additional capabilities described below. Two scenarios should be considered.

- 1) In the first scenario the client requests a network media channel, where all network media channels, the optical source and sink are compliant with the black link approach as specified in [ITU-T G.698.1] and [ITU-T G.698.2]. The request must include the end points of the network media channel being requested and the application code. Specific details on the network media channel center-frequency and spectral width may or may not be included in the request, depending on the control architecture of the implementation. The media controller must configure the non-associated overhead that is conveyed via the overhead communications channel (OCC) (defined in [ITU-T G.7712]), to support control/management of the network media channel. The client controller is responsible for ensuring that the OTSi provided across the interface is compliant with the network media channel requested and that the digital content of the optical signals passed across the network media channel can be successfully received at the far end. The media control function must also:
- Support the configuration of the requested media channel and the non-associated overhead.
 - Support the configuration and monitoring of the optical multiplex section (OMS) and optical transmission section (OTS) maintenance entities.

Where network media channels are not compliant with in [ITU-T G.698.1] or [ITU-T G.698.2], joint engineering is required to ensure proper interoperability between the client and media providers and their associated control functions (if separate). In this case, a connection request must include the end points of the network media channel being requested and a mutually agreed application identifier. Beyond that, the request contents are dependent on the decisions reached in the joint engineering of the network and control/management scheme employed.

² For example, a fibre may only be carrying one other signal currently, but it is intended to support up to a total of 80 signals

³ The implementation of a media control function is vendor specific; it may be offered as a stand-alone application or integrated with the control function for digital layer networks.

- 2) In the second scenario the client network requests an OTU network connection. In this case the media control function must:
- Accept requests for an OTU connection that is translated into an OTSiA connection.
 - This requires the capability to configure the media channels, the non-associated overhead and the optical termination source and sink. Some of the OTSiA OAM functions are supported by non-associated overhead and this must be configured when an OTSi is assigned to a network media channel.
 - As described above, ensure that the OTSi and network media channel are compatible.
 - Ensure that the OTU supported by the OTSi/OTSiA passed across the network media channel is delivered successfully to the destination interface.
 - Support the configuration of media channels to support the OTSiA.
 - Support configuration and monitoring of the OMS and OTS maintenance entities.

Bibliography

- [b-ITU-T Y.3300] Recommendation ITU-T Y.3300 (2014), *Framework of software-defined networking*.
- [b-UML] G. Booch, J. Rumbaugh and I Jacobson (2005), *The Unified Modelling Language Reference Manual* Second edition, Addison-Wesley.
- [b-ONF] Open Networking Foundation (2016), *SDN Architecture 1.1*, TR-521, February.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
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	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
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, and associated measurements and tests
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, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities
Series Z	Languages and general software aspects for telecommunication systems