

Network Working Group
Request for Comments: 4258
Category: Informational

D. Brungard, Ed.
ATT
November 2005

Requirements for Generalized Multi-Protocol Label Switching (GMPLS)
Routing for the Automatically Switched Optical Network (ASON)

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The Generalized Multi-Protocol Label Switching (GMPLS) suite of protocols has been defined to control different switching technologies as well as different applications. These include support for requesting Time Division Multiplexing (TDM) connections including Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) and Optical Transport Networks (OTNs).

This document concentrates on the routing requirements placed on the GMPLS suite of protocols in order to support the capabilities and functionalities of an Automatically Switched Optical Network (ASON) as defined by the ITU-T.

Table of Contents

1. Introduction	2
2. Conventions Used in This Document	4
3. ASON Routing Architecture and Requirements	4
3.1. Multiple Hierarchical Levels of ASON Routing Areas (RAs) ...	5
3.2. Hierarchical Routing Information Dissemination	6
3.3. Configuration	8
3.3.1. Configuring the Multi-Level Hierarchy	8
3.3.2. Configuring RC Adjacencies	8
3.4. Evolution	8
3.5. Routing Attributes	8
3.5.1. Taxonomy of Routing Attributes	9
3.5.2. Commonly Advertised Information	9
3.5.3. Node Attributes	10
3.5.4. Link Attributes	11
4. Security Considerations	12
5. Conclusions	12
6. Contributors	15
7. Acknowledgements	15
8. References	16
8.1. Normative References	16
8.2. Informative References	16

1. Introduction

The Generalized Multi-Protocol Label Switching (GMPLS) suite of protocols provides, among other capabilities, support for controlling different switching technologies. These include support for requesting TDM connections utilizing SONET/SDH (see [T1.105] and [G.707], respectively) as well as Optical Transport Networks (OTNs, see [G.709]). However, there are certain capabilities that are needed to support the ITU-T G.8080 control plane architecture for an Automatically Switched Optical Network (ASON). Therefore, it is desirable to understand the corresponding requirements for the GMPLS protocol suite. The ASON control plane architecture is defined in [G.8080]; ASON routing requirements are identified in [G.7715] and in [G.7715.1] for ASON link state protocols. These Recommendations apply to all [G.805] layer networks (e.g., SDH and OTN), and provide protocol-neutral functional requirements and architecture.

This document focuses on the routing requirements for the GMPLS suite of protocols to support the capabilities and functionality of ASON control planes. This document summarizes the ASON requirements using ASON terminology. This document does not address GMPLS applicability or GMPLS capabilities. Any protocol (in particular, routing)

applicability, design, or suggested extensions are strictly outside the scope of this document. ASON (Routing) terminology sections are provided in Appendixes 1 and 2.

The ASON routing architecture is based on the following assumptions:

- A network is subdivided based on operator decision and criteria (e.g., geography, administration, and/or technology); the network subdivisions are defined in ASON as Routing Areas (RAs).
- The routing architecture and protocols applied after the network is subdivided are an operator's choice. A multi-level hierarchy of RAs, as defined in ITU-T [G.7715] and [G.7715.1], provides for a hierarchical relationship of RAs based on containment; i.e., child RAs are always contained within a parent RA. The hierarchical containment relationship of RAs provides for routing information abstraction, thereby enabling scalable routing information representation. The maximum number of hierarchical RA levels to be supported is not specified (outside the scope of this document).
- Within an ASON RA and for each level of the routing hierarchy, multiple routing paradigms (hierarchical, step-by-step, source-based), centralized or distributed path computation, and multiple different routing protocols MAY be supported. The architecture does not assume a one-to-one correspondence between a routing protocol and an RA level, and allows the routing protocol(s) used within different RAs (including child and parent RAs) to be different. The realization of the routing paradigm(s) to support the hierarchical levels of RAs is not specified.
- The routing adjacency topology (i.e., the associated Protocol Controller (PC) connectivity) and transport topology are NOT assumed to be congruent.
- The requirements support architectural evolution, e.g., a change in the number of RA levels, as well as aggregation and segmentation of RAs.

The description of the ASON routing architecture provides for a conceptual reference architecture, with definition of functional components and common information elements to enable end-to-end routing in the case of protocol heterogeneity and facilitate management of ASON networks. This description is only conceptual: no physical partitioning of these functions is implied.

2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Although [RFC2119] describes interpretations of these key words in terms of protocol specifications and implementations, they are used in this document to describe design requirements for protocol extensions.

3. ASON Routing Architecture and Requirements

The fundamental architectural concept is the RA and its related functional components (see Appendix 2 on terminology). The routing services offered by an RA are provided by a Routing Performer (RP). An RP is responsible for a single RA, and it MAY be functionally realized using distributed Routing Controllers (RCs). The RC, itself, MAY be implemented as a cluster of distributed entities (ASON refers to the cluster as a Routing Control Domain (RCD)). The RC components for an RA receive routing topology information from their associated Link Resource Manager(s) (LRMs) and store this information in the Routing Information Database (RDB). The RDB is replicated at each RC bounded to the same RA, and MAY contain information about multiple transport plane network layers. Whenever the routing topology changes, the LRM informs the corresponding RC, which in turn updates its associated RDB. In order to ensure RDB synchronization, the RCs cooperate and exchange routing information. Path computation functions MAY exist in each RC, MAY exist on selected RCs within the same RA, or MAY be centralized for the RA.

In this context, communication between RCs within the same RA is realized using a particular routing protocol (or multiple protocols). In ASON, the communication component is represented by the protocol controller (PC) component(s) and the protocol messages are conveyed over the ASON control plane's Signaling Control Network (SCN). The PC MAY convey information for one or more transport network layers (refer to the note in Section 3.2). The RC is protocol independent, and RC communications MAY be realized by multiple, different PCs within an RA.

The ASON routing architecture defines a multi-level routing hierarchy of RAs based on a containment model to support routing information abstraction. [G.7715.1] defines the ASON hierarchical link state routing protocol requirements for communication of routing information within an RA (one level) to support hierarchical routing information dissemination (including summarized routing information

for other levels). The communication between any of the other functional component(s) (e.g., SCN, LRM, and between RCDs (RC-RC communication between RAs)) is outside the scope of [G.7715.1] protocol requirements and, thus, is also outside the scope of this document.

ASON routing components are identified by identifiers that are drawn from different name spaces (see [G.7715.1]). These are control plane identifiers for transport resources, components, and SCN addresses. The formats of those identifiers in a routing protocol realization SHALL be implementation specific and outside the scope of this document.

The failure of an RC, or the failure of communications between RCs, and the subsequent recovery from the failure condition MUST NOT disrupt calls in progress (i.e., already established) and their associated connections. Calls being set up MAY fail to complete, and the call setup service MAY be unavailable during recovery actions.

3.1. Multiple Hierarchical Levels of ASON Routing Areas (RAs)

[G.8080] introduces the concept of a Routing Area (RA) in reference to a network subdivision. RAs provide for routing information abstraction. Except for the single RA case, RAs are hierarchically contained: a higher-level (parent) RA contains lower-level (child) RAs that in turn MAY also contain RAs, etc. Thus, RAs contain RAs that recursively define successive hierarchical RA levels.

However, the RA containment relationship describes only an architectural hierarchical organization of RAs. It does not restrict a specific routing protocol's realization (e.g., OSPF multi-areas, path computation, etc.). Moreover, the realization of the routing paradigm to support a hierarchical organization of RAs and the number of hierarchical RA levels to be supported is routing protocol specific and outside the scope of this document.

In a multi-level hierarchy of RAs, it is necessary to distinguish among RCs for the different levels of the RA hierarchy. Before any pair of RCs establishes communication, they MUST verify that they are bound to the same parent RA (see Section 3.2). An RA identifier (RA ID) is required to provide the scope within which the RCs can communicate. To distinguish between RCs bound to the same RA, an RC identifier (RC ID) is required; the RC ID MUST be unique within its containing RA.

An RA represents a partition of the data plane, and its identifier (i.e., RA ID) is used within the control plane as a reference to the data plane partition. Each RA within a carrier's network SHALL be

uniquely identifiable. RA IDs MAY be associated with a transport plane name space, whereas RC IDs are associated with a control plane name space.

3.2. Hierarchical Routing Information Dissemination

Routing information can be exchanged between RCs bound to adjacent levels of the RA hierarchy, i.e., Level N+1 and N, where Level N represents the RAs contained by Level N+1. The links connecting RAs may be viewed as external links (inter-RA links), and the links representing connectivity within an RA may be viewed as internal links (intra-RA links). The external links to an RA at one level of the hierarchy may be internal links in the parent RA. Intra-RA links of a child RA MAY be hidden from the parent RA's view.

The physical location of RCs for adjacent RA levels, their relationship, and their communication protocol(s) are outside the scope of this document. No assumption is made regarding how RCs communicate between adjacent RA levels. If routing information is exchanged between an RC, its parent, and its child RCs, it SHOULD include reachability (see Section 3.5.3) and MAY include, upon policy decision, node and link topology. Communication between RAs only takes place between RCs with a parent/child relationship. RCs of one RA never communicate with RCs of another RA at the same level. There SHOULD not be any dependencies on the different routing protocols used within an RA or in different RAs.

Multiple RCs bound to the same RA MAY transform (filter, summarize, etc.) and then forward information to RCs at different levels. However, in this case, the resulting information at the receiving level must be self-consistent (i.e., ensure consistency between transform operations performed on routing information at different levels to ensure proper information processing). This MAY be achieved using a number of mechanisms.

Note: There is no implied relationship between multi-layer transport networks and multi-level routing. Implementations MAY support a hierarchical routing topology (multi-level) with a single routing protocol instance for multiple transport switching layers or a hierarchical routing topology for one transport switching layer.

1. Type of Information Exchanged

The type of information flowing upward (i.e., Level N to Level N+1) and the information flowing downward (i.e., Level N+1 to Level N) are used for similar purposes, namely, the exchange of reachability information and summarized topology information to

allow routing across multiple RAs. The summarization of topology information may impact the accuracy of routing and may require additional path calculation.

The following information exchanges are expected:

- Level N+1 visibility to Level N reachability and topology (or upward information communication) allowing RC(s) at Level N+1 to determine the reachable endpoints from Level N.
- Level N visibility to Level N+1 reachability and topology (or downward information communication) allowing RC(s) bounded to an RA at Level N to develop paths to reachable endpoints outside of the RA.

2. Interactions between Upward and Downward Communication

When both upward and downward information exchanges contain endpoint reachability information, a feedback loop could potentially be created. Consequently, the routing protocol **MUST** include a method to:

- prevent information propagated from a Level N+1 RA's RC into the Level N RA's RC from being re-introduced into the Level N+1 RA's RC, and
- prevent information propagated from a Level N-1 RA's RC into the Level N RA's RC from being re-introduced into the Level N-1 RA's RC.

The routing protocol **SHALL** differentiate the routing information originated at a given-level RA from derived routing information (received from external RAs), even when this information is forwarded by another RC at the same level. This is a necessary condition to be fulfilled by routing protocols to be loop free.

3. Method of Communication

Two approaches exist for communication between Level N and N+1:

- The first approach places an instance of a Level N routing function and an instance of a Level N+1 routing function in the same system. The communications interface is within a single system and is thus not an open interface subject to standardization. However, information re-advertisement or leaking **MUST** be performed in a consistent manner to ensure interoperability and basic routing protocol correctness (e.g., cost/metric value).

- The second approach places the Level N routing function on a separate system from the Level N+1 routing function. In this case, a communication interface must be used between the systems containing the routing functions for different levels. This communication interface and mechanisms are outside the scope of this document.

3.3. Configuration

3.3.1. Configuring the Multi-Level Hierarchy

The RC MUST support static (i.e., operator assisted) and MAY support automated configuration of the information describing its relationship to its parent and its child within the hierarchical structure (including RA ID and RC ID). When applied recursively, the whole hierarchy is thus configured.

3.3.2. Configuring RC Adjacencies

The RC MUST support static (i.e., operator assisted) and MAY support automated configuration of the information describing its associated adjacencies to other RCs within an RA. The routing protocol SHOULD support all the types of RC adjacencies described in Section 9 of [G.7715]. The latter includes congruent topology (with distributed RC) and hubbed topology (e.g., note that the latter does not automatically imply a designated RC).

3.4. Evolution

The containment relationships of RAs may change, motivated by events such as mergers, acquisitions, and divestitures.

The routing protocol SHOULD be capable of supporting architectural evolution in terms of the number of hierarchical levels of RAs, as well as the aggregation and segmentation of RAs. RA ID uniqueness within an administrative domain may facilitate these operations. The routing protocol is not expected to automatically initiate and/or execute these operations. Reconfiguration of the RA hierarchy may not disrupt calls in progress, though calls being set up may fail to complete, and the call setup service may be unavailable during reconfiguration actions.

3.5. Routing Attributes

Routing for transport networks is performed on a per-layer basis, where the routing paradigms MAY differ among layers and within a layer. Not all equipment supports the same set of transport layers or the same degree of connection flexibility at any given layer. A

server layer trail may support various clients, involving different adaptation functions. In addition, equipment may support variable adaptation functionality, whereby a single server layer trail dynamically supports different multiplexing structures. As a result, routing information MAY include layer-specific, layer-independent, and client/server adaptation information.

3.5.1. Taxonomy of Routing Attributes

Attributes can be organized according to the following categories:

- Node related or link related
- Provisioned, negotiated, or automatically configured
- Inherited or layer specific (client layers can inherit some attributes from the server layer, while other attributes such as Link Capacity are specified by layer)

(Component) link attributes MAY be statically or automatically configured for each transport network layer. This may lead to unnecessary repetition. Hence, the inheritance property of attributes MAY also be used to optimize the configuration process.

ASON uses the term SubNetwork Point (SNP) for the control plane representation of a transport plane resource. The control plane representation and transport plane topology are NOT assumed to be congruent; the control plane representation SHALL not be restricted by the physical topology. The relational grouping of SNPs for routing is termed an SNP Pool (SNPP). The routing function understands topology in terms of SNPP links. Grouping MAY be based on different link attributes (e.g., SRLG information, link weight, etc).

Two RAs may be linked by one or more SNPP links. Multiple SNPP links may be required when component links are not equivalent for routing purposes with respect to the RAs to which they are attached, to the containing RA, or when smaller groupings are required.

3.5.2. Commonly Advertised Information

Advertisements MAY contain the following common set of information regardless of whether they are link or node related:

- RA ID of the RA to which the advertisement is bounded
- RC ID of the entity generating the advertisement

- Information to uniquely identify advertisements
- Information to determine whether an advertisement has been updated
- Information to indicate when an advertisement has been derived from a different level RA

3.5.3. Node Attributes

All nodes belong to an RA; hence, the RA ID can be considered an attribute of all nodes. Given that no distinction is made between abstract nodes and those that cannot be decomposed any further, the same attributes MAY be used for their advertisement. In the following tables, Capability refers to the level of support required in the realization of a link state routing protocol, whereas Usage refers to the degree of operational control that SHOULD be available to the operator.

The following Node Attributes are defined:

Attribute -----	Capability -----	Usage -----
Node ID	REQUIRED	REQUIRED
Reachability	REQUIRED	OPTIONAL

Table 1. Node Attributes

Reachability information describes the set of endpoints that are reachable by the associated node. It MAY be advertised as a set of associated external (e.g., User Network Interface (UNI)) address/address prefixes or a set of associated SNPP link IDs/SNPP ID prefixes, the selection of which MUST be consistent within the applicable scope. These are control plane identifiers; the formats of these identifiers in a protocol realization are implementation specific and outside the scope of this document.

Note: No distinction is made between nodes that may have further internal details (i.e., abstract nodes) and those that cannot be decomposed any further. Hence, the attributes of a node are not considered as only single-switch attributes but MAY apply to a node at a higher level of the hierarchy that represents a subnetwork.

3.5.4. Link Attributes

The following Link Attributes are defined:

Link Attribute -----	Capability -----	Usage -----
Local SNPP link ID	REQUIRED	REQUIRED
Remote SNPP link ID	REQUIRED	REQUIRED
Layer Specific Characteristics	see Table 3	

Table 2. Link Attributes

The SNPP link ID MUST be sufficient to uniquely identify (within the Node ID scope) the corresponding transport plane resource, taking into account the separation of data and control planes (see Section 3.5.1; the control plane representation and transport plane topology are not assumed to be congruent). The SNPP link ID format is routing protocol specific.

Note: When the remote end of an SNPP link is located outside of the RA, the remote SNPP link ID is OPTIONAL.

The following link characteristic attributes are defined:

- Signal Type: This identifies the characteristic information of the layer network.
- Link Weight: This is the metric indicating the relative desirability of a particular link over another, e.g., during path computation.
- Resource Class: This corresponds to the set of administrative groups assigned by the operator to this link. A link MAY belong to zero, one, or more administrative groups.
- Local Connection Types: This attribute identifies whether the local SNP represents a Termination Connection Point (CP), a Connection Point (CP), or can be flexibly configured as a TCP.
- Link Capacity: This provides the sum of the available and potential bandwidth capacity for a particular network transport layer. Other capacity measures MAY be further considered.
- Link Availability: This represents the survivability capability such as the protection type associated with the link.
- Diversity Support: This represents diversity information such as the SRLG information associated with the link.

- **Local Adaptation Support:** This indicates the set of client layer adaptations supported by the TCP associated with the local SNPP. This is applicable only when the local SNP represents a TCP or can be flexibly configured as a TCP.

Link Characteristics	Capability	Usage
-----	-----	-----
Signal Type	REQUIRED	OPTIONAL
Link Weight	REQUIRED	OPTIONAL
Resource Class	REQUIRED	OPTIONAL
Local Connection Types	REQUIRED	OPTIONAL
Link Capacity	REQUIRED	OPTIONAL
Link Availability	OPTIONAL	OPTIONAL
Diversity Support	OPTIONAL	OPTIONAL
Local Adaptation Support	OPTIONAL	OPTIONAL

Table 3. Link Characteristics

Note: Separate advertisements of layer-specific attributes MAY be chosen. However, this may lead to unnecessary duplication. This can be avoided using the inheritance property, so that the attributes derivable from the local adaptation information do not need to be advertised. Thus, an optimization MAY be used when several layers are present by indicating when an attribute is inheritable from a server layer.

4. Security Considerations

The ASON routing protocol MUST deliver the operational security objectives where required. The overall security objectives (defined in ITU-T Recommendation [M.3016]) of confidentiality, integrity, and accountability may take on varying levels of importance. These objectives do not necessarily imply requirements on the routing protocol itself, and MAY be met by other established means.

Note: A threat analysis of a proposed routing protocol SHOULD address masquerade, eavesdropping, unauthorized access, loss or corruption of information (including replay attacks), repudiation, forgery, and denial of service attacks.

5. Conclusions

The description of the ASON routing architecture and components is provided in terms of routing functionality. This description is only conceptual: no physical partitioning of these functions is implied.

In summary, the ASON routing architecture assumes:

- A network is subdivided into ASON RAs, which MAY support multiple routing protocols; no one-to-one relationship SHALL be assumed.
- Routing Controllers (RCs) provide for the exchange of routing information (primitives) for the RA. The RC is protocol independent and MAY be realized by multiple, different protocol controllers within an RA. The routing information exchanged between RCs SHALL be subject to policy constraints imposed at reference points (External- and Internal-NNI).
- In a multi-level RA hierarchy based on containment, communication between RCs of different RAs happens only when there is a parent/child relationship between the RAs. RCs of child RAs never communicate with the RCs of other child RAs. There SHOULD not be any dependencies on the different routing protocols used within a child RA and that of its parent. The routing information exchanged within the parent RA SHALL be independent of both the routing protocol operating within a child RA and any control distribution choice(s), e.g., centralized, fully distributed.
- For an RA, the set of RCs is referred to as an ASON routing (control) domain. The routing information exchanged between routing domains (inter-RA, i.e., inter-domain) SHALL be independent of both the intra-domain routing protocol(s) and the intra-domain control distribution choice(s), e.g., centralized, fully distributed. RCs bounded to different RA levels MAY be collocated within the same physical element or physically distributed.
- The routing adjacency topology (i.e., the associated PC connectivity topology) and the transport network topology SHALL NOT be assumed to be congruent.
- The routing topology SHALL support multiple links between nodes and RAs.

In summary, the following functionality is expected from GMPLS routing to instantiate the ASON hierarchical routing architecture realization (see [G.7715] and [G.7715.1]):

- RAs SHALL be uniquely identifiable within a carrier's network, each having a unique RA ID within the carrier's network.
- Within an RA (one level), the routing protocol SHALL support dissemination of hierarchical routing information (including summarized routing information for other levels) in support of an

architecture of multiple hierarchical levels of RAs; the number of hierarchical RA levels to be supported by a routing protocol is implementation specific.

- The routing protocol SHALL support routing information based on a common set of information elements as defined in [G.7715] and [G.7715.1], divided between attributes pertaining to links and abstract nodes (each representing either a subnetwork or simply a node). [G.7715] recognizes that the manner in which the routing information is represented and exchanged will vary with the routing protocol used.
- The routing protocol SHALL converge such that the distributed RDBs become synchronized after a period of time.

To support hierarchical routing information dissemination within an RA, the routing protocol MUST deliver:

- Processing of routing information exchanged between adjacent levels of the hierarchy (i.e., Level N+1 and N) including reachability and, upon policy, decision summarized topology information.
- Self-consistent information at the receiving level resulting from any transformation (filter, summarize, etc.) and forwarding of information from one RC to RC(s) at different levels when multiple RCs are bound to a single RA.
- A mechanism to prevent the re-introduction of information propagated into the Level N RA's RC back to the adjacent level RA's RC from which this information has been initially received.

In order to support operator-assisted changes in the containment relationships of RAs, the routing protocol SHALL support evolution in terms of the number of hierarchical levels of RAs. For example: support of non-disruptive operations such as adding and removing RAs at the top/bottom of the hierarchy, adding or removing a hierarchical level of RAs in or from the middle of the hierarchy, as well as aggregation and segmentation of RAs. The number of hierarchical levels to be supported is routing protocol specific and reflects a containment relationship; e.g., an RA insertion involves supporting a different routing protocol domain in a portion of the network.

Reachability information (see Section 3.5.3) of the set of endpoints reachable by a node may be advertised either as a set of UNI Transport Resource addresses/address prefixes or a set of associated SNPP link IDs/SNPP link ID prefixes, assigned and selected consistently in their applicability scope. The formats of the

control plane identifiers in a protocol realization are implementation specific. Use of a routing protocol within an RA should not restrict the choice of routing protocols for use in other RAs (child or parent).

As ASON does not restrict the control plane architecture choice used, either a collocated architecture or a physically separated architecture may be used. A collection of links and nodes such as a subnetwork or RA MUST be able to represent itself to the wider network as a single logical entity with only its external links visible to the topology database.

6. Contributors

This document is the result of the CCAMP Working Group ASON Routing Requirements design team joint effort. The following are the design team member authors who contributed to the present document:

Wesam Alanqar (Sprint)
Deborah Brungard (ATT)
David Meyer (Cisco Systems)
Lyndon Ong (Ciena)
Dimitri Papadimitriou (Alcatel)
Jonathan Sadler (Tellabs)
Stephen Shew (Nortel)

7. Acknowledgements

The authors would like to thank Kireeti Kompella for having initiated the proposal of an ASON Routing Requirement Design Team and the ITU-T SG15/Q14 for their careful review and input.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

For information on the availability of the following documents, please see <http://www.itu.int>:

- [G.707] ITU-T Rec. G.707/Y.1322, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)", December 2003.
- [G.709] ITU-T Rec. G.709/Y.1331, "Interfaces for the Optical Transport Network (OTN)", March 2003.
- [G.7715] ITU-T Rec. G.7715/Y.1306, "Architecture and Requirements for the Automatically Switched Optical Network (ASON)", June 2002.
- [G.7715.1] ITU-T Draft Rec. G.7715.1/Y.1706.1, "ASON Routing Architecture and Requirements for Link State Protocols", November 2003.
- [G.805] ITU-T Rec. G.805, "Generic Functional Architecture of Transport Networks", March 2000.
- [G.8080] ITU-T Rec. G.8080/Y.1304, "Architecture for the Automatically Switched Optical Network (ASON)", November 2001 (and Revision, January 2003).
- [M.3016] ITU-T Rec. M.3016.0, "Security for the Management Plane: Overview", May 2005.
- [T1.105] ANSI T1.105, "Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates, and Formats", 2001.

Appendix 1: ASON Terminology

This document makes use of the following terms:

Administrative domain (see Recommendation [G.805]): For the purposes of [G.7715.1], an administrative domain represents the extent of resources that belong to a single player such as a network operator, a service provider, or an end-user. Administrative domains of different players do not overlap amongst themselves.

Adaptation function (see Recommendation [G.805]): A "transport processing function" that processes the client layer information for transfer over a server layer trail.

Client/Server relationship: The association between layer networks that is performed by an "adaptation" function to allow the link connection in the client layer network to be supported by a trail in the server layer network.

Control plane: Performs the call control and connection control functions. Through signaling, the control plane sets up and releases connections and may restore a connection in case of a failure.

(Control) Domain: Represents a collection of (control) entities that are grouped for a particular purpose. The control plane is subdivided into domains matching administrative domains. Within an administrative domain, further subdivisions of the control plane are recursively applied. A routing control domain is an abstract entity that hides the details of the RC distribution.

External NNI (E-NNI): Interfaces are located between protocol controllers between control domains.

Internal NNI (I-NNI): Interfaces are located between protocol controllers within control domains.

Link (see Recommendation [G.805]): A "topological component" that describes a fixed relationship between a "subnetwork" or "access group" and another "subnetwork" or "access group". Links are not limited to being provided by a single server trail.

Management plane: Performs management functions for the transport plane, the control plane, and the system as a whole. It also provides coordination between all the planes. The following management functional areas are performed in the management plane: performance, fault, configuration, accounting, and security management.

Management domain (see Recommendation [G.805]): A management domain defines a collection of managed objects that are grouped to meet organizational requirements according to geography, technology, policy, or other structure, and for a number of functional areas such as configuration, security, (FCAPS), for the purpose of providing control in a consistent manner. Management domains can be disjoint, contained, or overlapping. As such, the resources within an administrative domain can be distributed into several possible overlapping management domains. The same resource can therefore belong to several management domains simultaneously, but a management domain shall not cross the border of an administrative domain.

Multiplexing (see Recommendation [G.805]): Multiplexing techniques are used to combine client layer signals. The many-to-one relationship represents the case of several link connections of client layer networks supported by one server layer trail at the same time.

Subnetwork Point (SNP): The SNP is a control plane abstraction that represents an actual or potential transport plane resource. SNPs (in different subnetwork partitions) may represent the same transport resource. A one-to-one correspondence should not be assumed.

Subnetwork Point Pool (SNPP): A set of SNPs that are grouped together for the purposes of routing.

Termination Connection Point (TCP): A TCP represents the output of a Trail Termination function or the input to a Trail Termination Sink function.

Trail (see Recommendation [G.805]): A "transport entity" that consists of an associated pair of "unidirectional trails" capable of simultaneously transferring information in opposite directions between their respective inputs and outputs.

Transport plane: Provides bi-directional or unidirectional transfer of user information, from one location to another. It can also provide transfer of some control and network management information. The transport plane is layered; it is equivalent to the Transport Network defined in the [G.805] Recommendation.

User Network Interface (UNI): Interfaces are located between protocol controllers between a user and a control domain. Note: there is no routing function associated with a UNI reference point.

Variable adaptation function: A single server layer trail may dynamically support different multiplexing structures, i.e., link connections for multiple client layer networks.

Appendix 2: ASON Routing Terminology

This document makes use of the following terms:

Routing Area (RA): An RA represents a partition of the data plane, and its identifier is used within the control plane as the representation of this partition. Per [G.8080], an RA is defined by a set of subnetworks, the links that interconnect them, and the interfaces representing the ends of the links exiting that RA. An RA may contain smaller RAs inter-connected by links. The limit of subdivision results in an RA that contains two subnetworks interconnected by a single link.

Routing Database (RDB): Repository for the local topology, network topology, reachability, and other routing information that is updated as part of the routing information exchange and may additionally contain information that is configured. The RDB may contain routing information for more than one Routing Area (RA).

Routing Components: ASON routing architecture functions. These functions can be classified as protocol independent (Link Resource Manager or LRM, Routing Controller or RC) and protocol specific (Protocol Controller or PC).

Routing Controller (RC): Handles (abstract) information needed for routing and the routing information exchange with peering RCs by operating on the RDB. The RC has access to a view of the RDB. The RC is protocol independent.

Note: Since the RDB may contain routing information pertaining to multiple RAs (and possibly to multiple layer networks), the RCs accessing the RDB may share the routing information.

Link Resource Manager (LRM): Supplies all the relevant component and Traffic Engineering (TE) link information to the RC. It informs the RC about any state changes of the link resources it controls.

Protocol Controller (PC): Handles protocol-specific message exchanges according to the reference point over which the information is exchanged (e.g., E-NNI, I-NNI), and internal exchanges with the RC. The PC function is protocol dependent.

Authors' Addresses

Wesam Alanqar
Sprint

EMail: wesam.alanqar@mail.sprint.com

Deborah Brungard, Ed.
AT&T
Rm. D1-3C22 - 200 S. Laurel Ave.
Middletown, NJ 07748, USA

Phone: +1 732 4201573
EMail: dbrungard@att.com

David Meyer
Cisco Systems

EMail: dmm@1-4-5.net

Lyndon Ong
Ciena Corporation
5965 Silver Creek Valley Rd,
San Jose, CA 95128, USA

Phone: +1 408 8347894
EMail: lyong@ciena.com

Dimitri Papadimitriou
Alcatel
Francis Wellensplein 1,
B-2018 Antwerpen, Belgium

Phone: +32 3 2408491
EMail: dimitri.papadimitriou@alcatel.be

Jonathan Sadler
1415 W. Diehl Rd
Naperville, IL 60563

EMail: jonathan.sadler@tellabs.com

Stephen Shew
Nortel Networks
PO Box 3511 Station C
Ottawa, Ontario, CANADA K1Y 4H7

Phone: +1 613 7632462
EMail: sdsheew@nortelnetworks.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

