

CCAMP Working Group
Internet Draft
Category: Informational

Expiration Date: May 2004

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

December 2003

**Requirements for Generalized MPLS (GMPLS) Routing
for Automatically Switched Optical Network (ASON)**

[draft-ietf-ccamp-gmpls-ason-routing-reqts-01.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC-2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

The Generalized MPLS (GMPLS) suite of protocols has been defined to control different switching technologies as well as different applications. These include support for requesting TDM connections including SONET/SDH and Optical Transport Networks (OTNs).

This document concentrates on the routing requirements on the GMPLS suite of protocols to support the capabilities and functionalities of an Automatically Switched Optical Network (ASON).

1. Contributors

This document is the result of the CCAMP Working Group ASON Routing Requirements design team joint effort.

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](#).

3. Introduction

The GMPLS suite of protocol provides support for controlling different switching technologies as well as different applications. These include support for requesting TDM connections including SONET/SDH (see ANSI T1.105/ITU-T G.707) as well as Optical Transport Networks (see ITU-T G.709). However, there are certain capabilities that are needed to support the ITU-T G.8080 control plane architecture for the 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 [\[G.8080\]](#) and ASON routing requirements are identified in [\[G.7715\]](#) and refined in [\[G.7715.1\]](#) for link state architectures. These recommendations provide functional requirements and architecture, they provide a protocol neutral approach.

This document focuses on the routing requirements for the GMPLS suite of protocols to support the capabilities and functionalities of ASON control planes. It discusses the requirements for GMPLS routing that MAY subsequently lead to additional backward compatible extensions to support the capabilities specified in the above referenced document. A description of backward compatibility considerations is provided in [Section 5](#). Nonetheless, any protocol (in particular, routing) design or suggested protocol extensions is strictly outside the scope of this document. An ASON (Routing) terminology section is provided in Appendix 1 and Appendix 2.

The ASON model distinguishes reference points (representing points of protocol information exchange) defined (1) between an administrative domain and a user (user-network interface or UNI), (2) between administrative domains or within an administrative domain between different control domains (external network-network interface or E-NNI) and, (3) within the same administrative domain

between control components (or simply controllers) of the same control domain (internal network-network interface or I-NNI). The ASON model allows for the protocols used within different control domains to be different; and for the protocol used between control domains to be different than the protocols used within control domains. I-NNI interfaces are located between protocol controllers within a control domain. E-NNI interfaces are located on protocol controllers between control domains.

The term routing information refers to the abstract representation of network routing related information such as node and link attributes (see [Section 4.5](#)). No routing information is passed over the UNI. Routing information exchanged over the NNI is subject to the policy constraints at individual NNIs. The routing information exchanged over the E-NNI encapsulates the common semantics of the individual domain information while allowing different representation within each domain.

The ASON routing architecture is based on the following assumptions:

- A carrier's network is subdivided as Routing Areas (RAs). Each RA shall be uniquely identifiable within a carrier's network (i.e. administrative domain). Partitioning into RAs provides for routing information abstraction, thereby enabling scalable routing.
- Routing Controllers (RC) provide for the exchange of routing information between and within a RA. The routing information exchanged between RCs is subject to policy constraints imposed at reference points (E-NNI and I-NNI).
- A RA MAY support different routing protocols. There SHOULD NOT be any dependencies on the different routing protocols used.
- For a RA, the cluster of RCs is referred to as a routing domain. The routing information exchanged between routing domains (i.e. inter-domain) is independent of both the intra-domain routing protocol and the intra-domain control distribution choice(s), e.g. centralized, fully distributed.
- The routing adjacency topology and transport network topology SHALL NOT be assumed to be congruent.

The following functionality is expected from GMPLS routing to instantiate ASON routing realization (see [\[G.7715\]](#)):

- support multiple hierarchical levels of RAs
- support hierarchical routing information dissemination including summarized routing information
- support for multiple links between nodes and RAs (allowing for link and node diversity)

- support architectural evolution in terms of the number of levels of hierarchies, aggregation and segmentation of RAs
- 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 sub-network 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.

Also, the behaviour of GMPLS routing is expected to be such that:

- it is scalable with respect to the number of links, nodes and RAs
- in response to a routing event (e.g. topology update, reachability update), it delivers convergence and damping against flapping
- it fulfils the operational security objectives where required

4. ASON Requirements for GMPLS Routing

The description of the ASON routing components (see Appendix 2) is provided in terms of routing functionality. This description is only conceptual: no physical partitioning of these functions is implied.

The Routing Controller (RC) component receive routing information from their associated Link Resource Manager(s) (LRMs) regarding TE links and store this information in the Routing Information Database (RDB). The RDB is replicated at each RC within the same Routing Area (RA), and MAY contain information about multiple transport plane network layers. Whenever the state of a TE link (or component link) changes, the LRM informs the corresponding RC, which in turn updates its associated RDB. In order to assure RDB synchronization, the RCs co-operate and exchange routing information. In this context, communication between RCs is realized using a particular routing protocol represented by the protocol controller (PC) component and the protocol messages are conveyed over the Signaling Control Network (SCN). The PC MAY convey information for one or more transport network layers.

Note: the RC can be thought as the function processing the TE database populated by the link local/remote component and TE links (LRM) and by the network wide TE links through the PC which processes the protocol specific routing exchanges. The SCN corresponds to the IP control plane topology enabling routing exchanges between GMPLS controllers (i.e. the routing adjacencies).

The next sections detail the requirements for GMPLS routing to

support the following ASON routing functions.

4.1 Multiple Hierarchical Levels

Routing Areas (RAs) provide for routing information abstraction, thereby enabling scalable routing information representation. [G.7715] describes the use of hierarchy as one possible choice for routing area organization. RAs MAY be 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 routing levels. The realization of the routing paradigm to support hierarchical routing levels and the number of hierarchical levels to be supported is protocol specific and outside the scope of this document.

Note: an RA can be considered as representing either an Autonomous System (AS) or a canonical IGP routing area, both are sometimes referred to as routing regions (or simply regions).

ASON routing components are identified by values that MAY be drawn from several identifier spaces. The use of identifiers in a routing protocol realization is implementation specific and outside the scope of this document.

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

Note: RA IDs MAY be associated with a transport plane name space whereas RC IDs are associated with a control plane name space.

4.2 Hierarchical Routing Information Dissemination

Routing information MAY be exchanged between adjacent levels of the routing 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, and the links representing connectivity within an RA MAY be viewed as internal links.

The physical location of RCs at adjacent levels, their relationship and their communication protocol are outside the scope of this document. No assumption is made regarding how RCs communicate

between levels. Information exchange between an RC, its parent, and its child RCs, SHOULD include reachability and MAY include (upon policy decision) node and link topology.

Multiple RCs within a 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; this MAY be achieved using a number of mechanisms.

4.2.1 Communication between Adjacent Routing Levels

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 exchange 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) in 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 mechanism to prevent re-introduction of information propagated into the Level N RA back to the external level RA from which this information has been initially received.

The routing protocol is required to differentiate the routing information originated at a given level RA from the one derived using the routing information received from its external RAs (regardless of the level of the corresponding RCs). This is a necessary condition to be fulfilled by routing protocols to be loop free.

Also, for ASON, the routing information exchange may generate transient loops at the data plane if no route recording is used during signaling. So, at the data plane, it is not the routing exchange that guarantees (transient) loop avoidance but the signaling protocol by recording the route until the node where computation occurs (by excluding segments already traversed).

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.
- 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.

4.2.2 Configuring the Routing Hierarchy

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

4.2.3 Configuring RC Adjacencies

The RC MUST support static (i.e. operator assisted) and MAY support automated configuration of the information describing its control

adjacencies to other RCs within an RA. The protocol SHOULD support all the types of adjacencies described in Section 9 of [[G.7715](#)].

4.3 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 number of hierarchical levels, as well as aggregation and segmentation of RAs. RA IDs uniqueness within an administrative domain MAY facilitate these operations. The routing protocol is not expected to automatically initiate and/or execute these operations.

4.4 Multiple Links between Nodes and RAs

See [Section 4.5.3](#)

4.5 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 equipments support the same set of transport layers nor the same degree of connection flexibility at any given layer. A server layer trail may support various clients, involving different adaptation functions. Additionally, 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.

4.5.1 Taxonomy of 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 like Link Capacity are specified by layer).

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

TE links are configured through grouping of component 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 TE links. Multiple TE links may

be required when component links are not equivalent for routing purposes with respect to the RAs they are attached to, or to the containing RA, or when smaller groupings are required.

4.5.2 Commonly Advertised Information

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

- RA ID of 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 source external to the routing area

4.5.3 Node Attributes

All nodes belong to a 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.

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 address prefixes or a set of associated TE link IDs, consistently assigned within an administrative domain.

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.

4.5.4 Link Attributes

The following Link Attributes are defined:

Link Attribute	Capability	Usage
-----	-----	-----
Local TE link ID	REQUIRED	REQUIRED
Remote TE link ID	REQUIRED	REQUIRED
TE Link Characteristics	Table 3	

Table 2. Link Attributes

The TE link ID must be sufficient to uniquely identify the corresponding transport plane resource taking into account separation of data and control planes. The TE link ID format is routing protocol specific.

Note: when the remote end of a TE link is located outside of the RA, the remote TE link ID is OPTIONAL.

The following TE link characteristic attributes are defined:

- Signal Type: This identifies the characteristic information of the layer network.
- Link Weight: 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.
- Connection Types: This allows identification of whether the local component link is at a border or within an LSP region (see [[HIER](#)])
- 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 local component link associated to the local TE link. This can only exist when the "Local Connection Type" indicates crossing of an LSP Region or can be flexibly assigned to be at a border or within an LSP region (see [[HIER](#)]).

TE 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. TE 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 attributes derivable from the local adaptation information do not need to be advertised.

5. Backward Compatibility

Any particular realization of the ASON routing requirements MUST be backward compatible with the considered routing protocol(s).

Backward compatibility means that at any level of the routing hierarchy, nodes, some of which support the requirements described in this document, and some of which do not, must still be capable to operate as mandated by the OSPF, IS-IS, and/or IDR IETF WG and their corresponding GMPLS extensions (as mandated by the CCAMP IETF WG).

Additionally, nodes (that do not support these requirements) are made part of a multi-level routing hierarchy from their containing RA(s), must be capable of:

- rejecting any incoming routing information that would be addressed to them in a way that is not detrimental to the network as a whole
- communicating (at a given level) with any other node located at the same level and that implements these requirements

This assumes that such nodes do not communicate directly either with lower or upper level nodes.

6. Security Considerations

TBD.

7. Acknowledgements

The authors would like to thank Kireeti Kompella for having initiated the proposal of an ASON Routing Requirement Design Team.

8. Intellectual Property Considerations

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

W.Alanqar et al. - Expires May 2004

10

[draft-ietf-ccamp-gmpls-ason-routing-reqts-01.txt](#)

December 2003

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

9. References

- [RFC 2026] S.Bradner, "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [RFC 2119] S.Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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.8080] ITU-T Rec. G.8080/Y.1304, "Architecture for the Automatically Switched Optical Network (ASON)," November 2001 (and Revision, January 2003).
- [HIER] K.Kompella and Y.Rekhter, "LSP Hierarchy with Generalized MPLS TE," Internet draft (work in progress), [draft-ietf-mpls-lsp-hierarchy](#), Sept'02.

10. Author's Addresses

Wesam Alanqar (Sprint)
EMail: wesam.alanqar@mail.sprint.com

Deborah Brungard (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)

W.Alanqar et al. - Expires May 2004

11

[draft-ietf-ccamp-gmpls-ason-routing-regts-01.txt](#)

December 2003

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: sdshe@nortelnetworks.com

Appendix 1 - ASON Terminology

This document makes use of the following terms:

Administrative domain: See Recommendation G.805.

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 entities that are grouped for a particular purpose. G.8080 applies this G.805 recommendation concept (that defines two particular forms, the administrative domain and the management domain) to the control plane in the form of a control domain. The entities that are grouped in a control domain are components of the control plane.

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.

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.

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 G.805.

User Network Interface (UNI): interfaces are located between protocol controllers between a user and a control domain.

Appendix 2 - ASON Routing Terminology

This document makes use of the following terms:

Routing Area (RA): represents functionally either an Autonomous System (AS) or a canonical IGP routing area, both are sometimes referred to as routing regions (or simply regions).

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 hence possibly multiple layer networks), the RCs accessing the RDB may share the routing information.

- Link Resource Manager (LRM): supplies all the relevant component and 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.

Internal Links: links that are fully encapsulated by a routing area at a given level of hierarchy. Internal links to a child RA may be hidden from the parent RAs view.

External Links: links that are incident upon the routing area. Note that external links to a routing area at one level of the hierarchy may be internal links in the parent routing area.

Full Copyright Statement

"Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph

are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.