

Network Working Group
Request for Comments: 4397
Category: Informational

I. Bryskin
Independent Consultant
A. Farrel
Old Dog Consulting
February 2006

A Lexicography for the Interpretation of Generalized Multiprotocol
Label Switching (GMPLS) Terminology within the Context of the
ITU-T's Automatically Switched Optical Network (ASON) Architecture

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 (2006).

Abstract

Generalized Multiprotocol Label Switching (GMPLS) has been developed by the IETF to facilitate the establishment of Label Switched Paths (LSPs) in a variety of data plane technologies and across several architectural models. The ITU-T has specified an architecture for the control of Automatically Switched Optical Networks (ASON).

This document provides a lexicography for the interpretation of GMPLS terminology within the context of the ASON architecture.

It is important to note that GMPLS is applicable in a wider set of contexts than just ASON. The definitions presented in this document do not provide exclusive or complete interpretations of GMPLS concepts. This document simply allows the GMPLS terms to be applied within the ASON context.

Table of Contents

1. Introduction	3
2. Terminology	3
2.1. GMPLS Terminology Sources	3
2.2. ASON Terminology Sources	4
2.3. Common Terminology Sources	4
3. Lexicography	4
3.1. Network Presences	4
3.2. Resources	5
3.3. Layers	6
3.4. Labels	7
3.5. Data Links	7
3.6. Link Interfaces	8
3.7. Connections	9
3.8. Switching, Termination, and Adaptation Capabilities	10
3.9. TE Links and FAs	11
3.10. TE Domains	13
3.11. Component Links and Bundles	13
3.12. Regions	14
4. Guidance on the Application of this Lexicography	14
5. Management Considerations	15
6. Security Considerations	15
7. Acknowledgements	15
8. Normative References	16
9. Informative References	16

1. Introduction

Generalized Multiprotocol Label Switching (GMPLS) has been developed by the IETF to facilitate the establishment of Label Switched Paths (LSPs) in a variety of data plane technologies such as Packet Switching Capable (PSC), Layer Two Switching Capable (L2SC), Time Division Multiplexing (TDM), Lambda Switching Capable (LSC), and Fiber Switching Capable (FSC).

The ITU-T has specified an architecture for the control of Automatically Switched Optical Networks (ASON). This architecture forms the basis of many Recommendations within the ITU-T.

Because the GMPLS and ASON architectures were developed by different people in different standards bodies, and because the architectures have very different historic backgrounds (the Internet, and transport networks respectively), the terminology used is different.

This document provides a lexicography for the interpretation of GMPLS terminology within the context of the ASON architecture. This allows GMPLS documents to be generally understood by those familiar with ASON Recommendations. The definitions presented in this document do not provide exclusive or complete interpretations of the GMPLS concepts.

2. Terminology

Throughout this document, angle brackets (" $<$ " and " $>$ ") are used to indicate the context in which a term applies. For example, " $<$ Data Plane $>$ " as part of a description of a term means that the term applies within the data plane.

2.1. GMPLS Terminology Sources

GMPLS terminology is principally defined in [RFC3945]. Other documents provide further key definitions including [RFC4201], [RFC4202], [RFC4204], and [RFC4206].

The reader is recommended to become familiar with these other documents before attempting to use this document to provide a more general mapping between GMPLS and ASON.

For details of GMPLS signaling, please refer to [RFC3471] and [RFC3473]. For details of GMPLS routing, please refer to [RFC4203] and [RFC4205].

2.2. ASON Terminology Sources

The ASON architecture is specified in ITU-T Recommendation G.8080 [G-8080]. This is developed from generic functional architectures and requirements specified in [G-805], [G-807], and [G-872]. The ASON terminology is defined in several Recommendations in the ASON family such as [G-8080], [G-8081], [G-7713], [G-7714], and [G-7715]. The reader must be familiar with these documents before attempting to apply the lexicography set out in this document.

2.3. Common Terminology Sources

The work in this document builds on the shared view of ASON requirements and requirements expressed in [RFC4139], [RFC4258], and [RFC4394].

3. Lexicography

3.1. Network Presences

3.1.1. GMPLS Terms

Transport node <Data Plane> is a logical network device that is capable of originating and/or terminating of a data flow and/or switching it on the route to its destination.

Controller <Control Plane> is a logical entity that models all control plane intelligence (routing, traffic engineering (TE), and signaling protocols, path computation, etc.). A single controller can manage one or more transport nodes. Separate functions (such as routing and signaling) may be hosted at distinct sites and hence could be considered as separate logical entities referred to, for example, as the routing controller, the signaling controller, etc.

Label Switching Router (LSR) <Control & Data Planes> is a logical combination of a transport node and the controller that manages the transport node. Many implementations of LSRs collocate all control plane and data plane functions associated with a transport node within a single physical presence making the term LSR concrete rather than logical.

In some instances, the term LSR may be applied more loosely to indicate just the transport node or just the controller function dependent on the context.

Node <Control & Data Planes> is a synonym for an LSR.

Control plane network <Control Plane> is an IP network used for delivery of control plane (protocol) messages exchanged by controllers.

3.1.2. ASON Terms

A GMPLS transport node is an ASON network element.

A GMPLS controller is the set of ASON functional components controlling a given ASON network element (or partition of a network element). In ASON, this set of functional components may exist in one place or multiple places.

A GMPLS node is the combination of an ASON network element (or partition of a network element) and its associated control components.

The GMPLS control plane network is the ASON Signaling Communication Network (SCN). Note that both routing and signaling exchanges are carried by the SCN.

3.2. Resources

3.2.1. GMPLS Terms

Non-packet-based resource <Data Plane> is a channel of a certain bandwidth that could be allocated in a network data plane of a particular technology for the purpose of user traffic delivery. Examples of non-packet-based resources are timeslots, lambda channels, etc.

Packet-based resource <Data Plane> is an abstraction hiding the means related to the delivery of traffic with particular parameters (most importantly, bandwidth) with particular quality of service (QoS) over PSC media. Examples of packet-based resources are forwarding queues, schedulers, etc.

Layer Resource (Resource) <Data Plane>. A non-packet-based data plane technology may yield resources in different network layers (see section 3.3). For example, some TDM devices can operate with VC-12 timeslots, some with VC-4 timeslots, and some with VC4-4c timeslots. There are also multiple layers of packet-based resources (i.e., one per label in the label stack). Therefore, we define layer resource (or simply resource) irrespective of the underlying data plane technology as a basic data plane construct. It is defined by a combination of a particular data encoding type

and a switching/terminating bandwidth granularity. Examples of layer resources are: PSC1, PSC4, ATM VP, ATM VC, Ethernet, VC-12, VC-4, Lambda 10G, and Lambda 40G.

These three definitions give rise to the concept of Resource Type. Although not a formal term, this is useful shorthand to identify how and where a resource can be used dependent on the switching type, data encoding type, and switching/terminating bandwidth granularity (see section 3.8).

All other descriptions provided in this memo are tightly bound to the resource.

3.2.2. ASON Terms

ASON terms for resource:

- In the context of link discovery and resource management (allocation, binding into cross-connects, etc.), a GMPLS resource is one end of a link connection.
- In the context of routing, path computation, and signaling, a GMPLS resource is a link connection or trail termination.

Resource type is identified by a client CI (Characteristics Information) that could be carried by the resource.

3.3. Layers

3.3.1. GMPLS Terms

Layer <Data Plane> is a set of resources of the same type that could be used for establishing a connection or used for connectionless data delivery.

Note. In GMPLS, the existence of non-blocking switching function in a transport node in a particular layer is modeled explicitly as one of the functions of the link interfaces connecting the transport node to its data links.

A GMPLS layer is not the same as a GMPLS region. See section 3.12.

3.3.2. ASON Terms

A GMPLS layer is an ASON layer network.

3.4. Labels

3.4.1. GMPLS Terms

Label <Control Plane> is an abstraction that provides an identifier for use in the control plane in order to identify a transport plane resource.

3.4.2. ASON Terms

A GMPLS label is the portion of an ASON SNP name that follows the SNPP name.

3.5. Data Links

3.5.1. GMPLS Terms

Unidirectional data link end <Data Plane> is a set of resources that belong to the same layer and that could be allocated for the transfer of traffic in that layer from a particular transport node to the same neighboring transport node in the same direction. A unidirectional data link end is connected to a transport node by one or more link interfaces (see section 3.6).

Bidirectional data link end <Data Plane> is an association of two unidirectional data link ends that exist in the same layer and that could be used for the transfer of traffic in that layer between a particular transport node and the same neighbor in both directions. A bidirectional data link end is connected to a transport node by one or more link interfaces (see section 3.6).

Unidirectional data link <Data Plane> is an association of two unidirectional data link ends that exist in the same layer, that are connected to two transport nodes adjacent in that layer, and that could be used for the transfer of traffic between the two transport nodes in one direction.

Bidirectional data link <Data Plane> is an association of two bidirectional data link ends that exist in the same layer, that are connected to two transport nodes adjacent in that layer, and that could be used for the transfer of traffic between the two transport nodes in both directions.

3.5.2. ASON Terms

A GMPLS unidirectional data link end is a collection of connection points from the same client layer that are supported by a single trail termination (access point).

A GMPLS data link is an ASON link supported by a single server trail.

3.6. Link Interfaces

3.6.1. GMPLS Terms

Unidirectional link interface <Data Plane> is an abstraction that connects a transport node to a unidirectional data link end and represents (hides) the data plane intelligence like switching, termination, and adaptation in one direction. In GMPLS, link interfaces are often referred to as "GMPLS interfaces" and it should be understood that these are data plane interfaces and the term does not refer to the ability of a control plane interface to handle GMPLS protocols.

A single unidirectional data link end could be connected to a transport node by multiple link interfaces with one of them, for example, realizing switching function, while others realize the function of termination/adaptation.

Bidirectional link interface <Data Plane> is an association of two or more unidirectional link interfaces that connects a transport node to a bidirectional data link end and represents the data plane intelligence like switching, termination, and adaptation in both directions.

Link interface type <Data Plane> is identified by the function the interface provides. There are three distinct functions -- switching, termination, and adaptation; hence, there are three types of link interface. Thus, when a Wavelength Division Multiplexing (WDM) link can do switching for some lambda channels, and termination and TDM OC48 adaptation for some other lambda channels, we say that the link is connected to the transport node by three interfaces each of a separate type: switching, termination, and adaptation.

3.6.2. ASON Terms

A GMPLS interface is the set of trail termination and adaptation functions between one or more server layer trails and a specific client layer subnetwork (which commonly is a matrix in a network element).

The GMPLS interface type may be identified by the ASON adapted client layer, or by the terminated server layer, or a combination of the two, depending on the context. In some cases, a GMPLS interface comprises a set of ASON trail termination/adaptation functions, for which some connection points are bound to trail terminations and others to matrices.

3.7. Connections

3.7.1. GMPLS Terms

In GMPLS a connection is known as a Label Switched Path (LSP).

Unidirectional LSP (connection) <Data Plane> is a single resource or a set of cross-connected resources of a particular layer that could deliver traffic in that layer between a pair of transport nodes in one direction.

Unidirectional LSP (connection) <Control Plane> is the signaling state necessary to maintain a unidirectional data plane LSP.

Bidirectional LSP (connection) <Data Plane> is an association of two unidirectional LSPs (connections) that could simultaneously deliver traffic in a particular layer between a pair of transport nodes in opposite directions.

In the context of GMPLS, both unidirectional constituents of a bidirectional LSP (connection) take identical paths in terms of data links, are provisioned concurrently, and require a single (shared) control state.

Bidirectional LSP (connection) <Control Plane> is the signaling state necessary to maintain a bidirectional data plane LSP.

LSP (connection) segment <Data Plane> is a single resource or a set of cross-connected resources that constitutes a segment of an LSP (connection).

3.7.2. ASON Terms

A GMPLS LSP (connection) is an ASON network connection.

A GMPLS LSP segment is an ASON serial compound link connection.

3.8. Switching, Termination, and Adaptation Capabilities

3.8.1. GMPLS Terms

Switching capability <Data Plane> is a property (and defines a type) of a link interface that connects a particular data link to a transport node. This property/type characterizes the interface's ability to cooperate with other link interfaces connecting data links within the same layer to the same transport node for the purpose of binding resources into cross-connects. Switching capability is advertised as an attribute of the TE link local end associated with the link interface.

Termination capability <Data Plane> is a property of a link interface that connects a particular data link to a transport node. This property characterizes the interface's ability to terminate connections within the layer that the data link belongs to.

Adaptation capability <Data Plane> is a property of a link interface that connects a particular data link to a transport node. This property characterizes the interface's ability to perform a nesting function -- to use a locally terminated connection that belongs to one layer as a data link for some other layer.

The need for advertisement of adaptation and termination capabilities within GMPLS has been recognized, and work is in progress to determine how these will be advertised. It is likely that they will be advertised as a single combined attribute, or as separate attributes of the TE link local end associated with the link interface.

3.8.2. ASON Terms

In ASON applications:

The GMPLS switching capability is a property of an ASON link end representing its association with a matrix.

The GMPLS termination capability is a property of an ASON link end representing potential binding to a termination point.

The GMPLS adaptation capability is a property of an ASON link end representing potential adaptation to/from a client layer network.

3.9. TE Links and FAs

3.9.1. GMPLS Terms

TE link end <Control Plane> is a grouping for the purpose of advertising and routing of resources of a particular layer.

Such a grouping allows for decoupling of path selection from resource assignment. Specifically, a path could be selected in a centralized way in terms of TE link ends, while the resource assignment (resource reservation and label allocation) could be performed in a distributed way during the connection setup. A TE link end may reflect zero, one or more data link ends in the data plane. A TE link end is associated with exactly one layer.

TE link <Control Plane> is a grouping of two TE link ends associated with two neighboring transport nodes in a particular layer.

In contrast to a data link, which provides network flexibility in a particular layer and, therefore, is a "real" topological element, a TE link is a logical routing element. For example, an LSP path is computed in terms of TE links (or more precisely, in terms of TE link ends), while the LSP is provisioned over (that is, resources are allocated from) data links.

Virtual TE link is a TE link associated with zero data links.

TE link end advertising <Control Plane>. A controller managing a particular transport node advertises local TE link ends. Any controller in the TE domain makes a TE link available for its local path computation if it receives consistent advertisements of both TE link ends. Strictly speaking, there is no such thing as TE link advertising -- only TE link end advertising. TE link end advertising may contain information about multiple switching capabilities. This, however, should not be interpreted as advertising of a multi-layer TE link end, but rather as joint advertisement of ends of multiple parallel TE links, each representing resources in a separate layer. The advertisement may contain attributes shared by all TE links in the group (for example, protection capabilities, Shared Risk Link Groups (SRLGs), etc.), separate information related to each TE link (for example, switching capability, data encoding, unreserved bandwidth, etc.) as well as information related to inter-layer relationships of the advertised resources (for example, termination and adaptation capabilities) should the control plane decide to use them as the termination points of higher-layer data links. These higher-layer data links, however, are not real yet -- they are abstract until the underlying connections are established in the lower layers.

LSPs created in lower layers for the purpose of providing data links (extra network flexibility) in higher layers are called hierarchical connections or LSPs (H-LSPs), or simply hierarchies. LSPs created for the purpose of providing data links in the same layer are called stitching segments. H-LSPs and stitching segments could, but do not have to, be advertised as TE links. Naturally, if they are advertised as TE links (LSPs advertised as TE links are often referred to as TE-LSPs), they are made available for path computations performed on any controller within the TE domain into which they are advertised. H-LSPs and stitching segments could be advertised either individually or in TE bundles. An H-LSP or a stitching segment could be advertised as a TE link either into the same or a separate TE domain compared to the one within which it was provisioned.

A set of H-LSPs that is created (or could be created) in a particular layer to provide network flexibility (data links) in other layers is called a Virtual Network Topology (VNT). A single H-LSP could provide several (more than one) data links (each in a different layer).

Forwarding Adjacency (FA) <Control Plane> is a TE link that does not require a direct routing adjacency (peering) between the controllers managing its ends in order to guarantee control plane connectivity (a control channel) between the controllers. An example of an FA is an H-LSP or stitching segment advertised as a TE link into the same TE domain within which it was dynamically provisioned. In such cases, the control plane connectivity between the controllers at the ends of the H-LSP/stitching segment is guaranteed by the concatenation of control channels interconnecting the ends of each of its constituents. In contrast, an H-LSP or stitching segment advertised as a TE link into a TE domain (different than one where it was provisioned) generally requires a direct routing adjacency to be established within the TE domain where the TE link is advertised in order to guarantee control plane connectivity between the TE link ends. Therefore, is not an FA.

3.9.2. ASON Terms

The ITU term for a TE link end is Subnetwork Point (SNP) pool (SNPP).

The ITU term for a TE link is SNPP link.

The ITU term for an H-LSP is trail.

3.10. TE Domains

3.10.1 GMPLS Terms

TE link attribute is a parameter of the set of resources associated with a TE link end that is significant in the context of path computation.

Full TE visibility is a situation when a controller receives all unmodified TE advertisements from every other controller in a particular set of controllers.

Limited TE visibility is a situation when a controller receives summarized TE information, or does not receive TE advertisements from at least one of a particular set of controllers.

TE domain is a set of controllers each of which has full TE visibility within the set.

TE database (TED) is a memory structure within a controller that contains all TE advertisements generated by all controllers within a particular TE domain.

Vertical network integration is a set of control plane mechanisms and coordinated data plane mechanisms that span multiple layers. The control plane mechanisms exist on one or more controllers and operate either within a single control plane instance or between control plane instances. The data plane mechanisms consist of collaboration and adaptation between layers within a single transport node.

Horizontal network integration is a set of control plane mechanisms and coordinated data plane mechanisms that span multiple TE domains within the same layer. The control plane mechanisms exist on one or more controllers and operate either within a single control plane instance or between control plane instances. The data plane mechanisms consist of collaboration between TE domains.

3.11. Component Links and Bundles

3.11.1. GMPLS Terms

Component link end <Control Plane> is a grouping of resources of a particular layer that is not advertised as an individual TE link end. A component link end could represent one or more data link ends or any subset of resources that belong to one or more data link ends.

Component link <Control Plane> is a grouping of two or more component link ends associated with neighboring transport nodes (that is, directly interconnected by one or more data links) in a particular layer. Component links are equivalent to TE links except that the component link ends are not advertised separately.

TE bundle <Control Plane> is an association of several parallel (that is, connecting the same pair of transport nodes) component links whose attributes are identical or whose differences are sufficiently negligible that the TE domain can view the entire association as a single TE link. A TE bundle is advertised in the same way as a TE link, that is, by representing the associated component link ends as a single TE link end (TE bundle end) which is advertised.

3.12. Regions

3.12.1. GMPLS Terms

TE region <Control Plane> is a set of one or more layers that are associated with the same type of data plane technology. A TE region is sometimes called an LSP region or just a region. Examples of regions are: IP, ATM, TDM, photonic, fiber switching, etc. Regions and region boundaries are significant for the signaling sub-system of the control plane because LSPs are signaled substantially differently (i.e., use different signaling object formats and semantics) in different regions. Furthermore, advertising, routing, and path computation could be performed differently in different regions. For example, computation of paths across photonic regions requires a wider set of constraints (e.g., optical impairments, wavelength continuity, etc) and needs to be performed in different terms (e.g., in terms of individual resources -- lambda channels, rather than in terms of TE links) compared to path computation in other regions like IP or TDM.

4. Guidance on the Application of this Lexicography

As discussed in the introduction to this document, this lexicography is intended to bring the concepts and terms associated with GMPLS into the context of the ITU-T's ASON architecture. Thus, it should help those familiar with ASON to see how they may use the features and functions of GMPLS in order to meet the requirements of an ASON. For example, service providers wishing to establish a protected end-to-end service might read [SEG-PROT] and [E2E-PROT] and wish to understand how the GMPLS terms used relate to the ASON architecture so that they can confirm that they will satisfy their requirements.

This lexicography should not be used in order to obtain or derive definitive definitions of GMPLS terms. To obtain definitions of GMPLS terms that are applicable across all GMPLS architectural models, the reader should refer to the RFCs listed in the references sections of this document. [RFC3945] provides an overview of the GMPLS architecture and should be read first.

5. Management Considerations

Both GMPLS and ASON networks require management. Both GMPLS and ASON specifications include considerable efforts to provide operator control and monitoring, as well as Operations and Management (OAM) functionality.

These concepts are, however, out of scope of this document.

6. Security Considerations

Security is also a significant requirement of both GMPLS and ASON architectures.

Again, however, this informational document is intended only to provide a lexicography, and the security concerns are, therefore, out of scope.

7. Acknowledgements

The authors would like to thank participants in the IETF's CCAMP working group and the ITU-T's Study Group 15 for their help in producing this document. In particular, all those who attended the Study Group 15 Question 14 Interim Meeting in Holmdel, New Jersey during January 2005. Further thanks to all participants of Study Group 15 Questions 12 and 14 who have provided valuable discussion, feedback and suggested text.

Many thanks to Ichiro Inoue for his useful review and input, and to Scott Brim and Dimitri Papadimitriou for lengthy and constructive discussions. Ben Mack-Crane and Jonathan Sadler provided very helpful reviews and discussions of ASON terms. Thanks to Deborah Brungard and Kohei Shiomoto for additional review comments.

8. Normative References

- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.
- [RFC4202] Kompella, K. and Y. Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC4204] Lang, J., Ed., "Link Management Protocol (LMP)", RFC 4204, October 2005.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.

9. Informative References

- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC4139] Papadimitriou, D., Drake, J., Ash, J., Farrel, A., and L. Ong, "Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON)", RFC 4139, July 2005.
- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.
- [RFC4205] Kompella, K., Ed. and Y. Rekhter, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4205, October 2005.

- [RFC4258] Brungard, D., Ed., "Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Routing for the Automatically Switched Optical Network (ASON)", RFC 4258, November 2005.
- [RFC4394] Fedyk, D., Aboul-Magd, O., Brungard, D., Lang, J., and D. Papadimitriou, "A Transport Network View of the Link Management Protocol (LMP)", RFC 4394, February 2006.
- [E2E-PROT] Lang, J., Ed., Rekhter, Y., Ed., and D. Papadimitriou, D., Ed., "RSVP-TE Extensions in support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery", Work in Progress, April 2005.
- [SEG-PROT] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Based Segment Recovery", Work in Progress, May 2005.

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

- [G-8080] ITU-T Recommendation G.8080/Y.1304, Architecture for the automatically switched optical network (ASON).
- [G-805] ITU-T Recommendation G.805 (2000), Generic functional architecture of transport networks.
- [G-807] ITU-T Recommendation G.807/Y.1302 (2001), Requirements for the automatic switched transport network (ASTN).
- [G-872] ITU-T Recommendation G.872 (2001), Architecture of optical transport networks.
- [G-8081] ITU-T Recommendation G.8081 (2004), Terms and definitions for Automatically Switched Optical Networks (ASON).
- [G-7713] ITU-T Recommendation G.7713 (2001), Distributed Call and Connection Management.
- [G-7714] ITU-T Recommendation G.7714 Revision (2005), Generalized automatic discovery techniques.

[G-7715] ITU-T Recommendation G.7715 (2002), Architecture and Requirements for the Automatically Switched Optical Network (ASON).

Authors' Addresses

Igor Bryskin
Independent Consultant

E-Mail: i_bryskin@yahoo.com

Adrian Farrel
Old Dog Consulting

Phone: +44 (0) 1978 860944
E-Mail: adrian@olddog.co.uk

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 provided by the IETF Administrative Support Activity (IASA).