

Network Working Group
Request for Comments: 4394
Category: Informational

D. Fedyk
O. Aboul-Magd
Nortel Networks
D. Brungard
AT&T
J. Lang
Sonos, Inc.
D. Papadimitriou
Alcatel
February 2006

A Transport Network View of the Link Management Protocol (LMP)

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

The Link Management Protocol (LMP) has been developed as part of the Generalized MPLS (GMPLS) protocol suite to manage Traffic Engineering (TE) resources and links. The GMPLS control plane (routing and signaling) uses TE links for establishing Label Switched Paths (LSPs). This memo describes the relationship of the LMP procedures to 'discovery' as defined in the International Telecommunication Union (ITU-T), and ongoing ITU-T work. This document provides an overview of LMP in the context of the ITU-T Automatically Switched Optical Networks (ASON) and transport network terminology and relates it to the ITU-T discovery work to promote a common understanding for progressing the work of IETF and ITU-T.

Table of Contents

- 1. Introduction2
- 2. ASON Terminology and Abbreviations Related to Discovery3
 - 2.1. Terminology3
 - 2.2. Abbreviations4
- 3. Transport Network Architecture5
 - 3.1. G.8080 Discovery Framework7
- 4. Discovery Technologies9
 - 4.1. Generalized Automatic Discovery Techniques G.77149
 - 4.2. LMP and G.8080 Terminology Mapping9
 - 4.2.1. TE Link Definition and Scope12
 - 4.3. LMP and G.8080 Discovery Relationship13
 - 4.4. Comparing LMP and G.808014
- 5. Security Considerations15
- 6. Informative References15
- 7. Acknowledgements16

1. Introduction

The GMPLS control plane consists of several building blocks as described in [RFC3945]. The building blocks include signaling, routing, and link management for establishing LSPs. For scalability purposes, multiple physical resources can be combined to form a single TE link for the purposes of path computation and GMPLS control plane signaling.

As manual provisioning and management of these links are impractical in large networks, LMP was specified to manage TE links. Two mandatory management capabilities of LMP are control channel management and TE link property correlation. Additional optional capabilities include verifying physical connectivity and fault management. [LMP] defines the messages and procedures for GMPLS TE link management. [LMP-TEST] defines SONET/SDH-specific messages and procedures for link verification.

ITU-T Recommendation G.8080 Amendment 1 [G.8080] defines control plane discovery as two separate processes; one process occurs within the transport plane space and the other process occurs within the control plane space.

The ITU-T has developed Recommendation G.7714, "Generalized automatic discovery techniques" [G.7714], defining the functional processes and information exchange related to transport plane discovery aspects, i.e., layer adjacency discovery and physical media adjacency discovery. Specific methods and protocols are not defined in Recommendation G.7714. ITU-T Recommendation G.7714.1, "Protocol for automatic discovery in SDH and OTN networks" [G.7714.1], defines a

protocol and procedure for transport plane layer adjacency discovery (e.g., discovering the transport plane layer endpoint relationships and verifying their connectivity). The ITU-T is currently working to extend discovery to control plane aspects providing detail on a discovery framework architecture in G.8080 and a new Recommendation on "Control plane initial establishment, reconfiguration".

2. ASON Terminology and Abbreviations Related to Discovery

ITU-T Recommendation G.8080 Amendment 1 [G.8080] and ITU-T Recommendation G.7714 [G.7714] provide definitions and mechanisms related to transport plane discovery.

Note that in the context of this work, "Transport" relates to the data plane (sometimes called the transport plane or the user plane) and does not refer to the transport layer (layer 4) of the OSI seven layer model, nor to the concept of transport intended by protocols such as the Transmission Control Protocol (TCP).

Special care must be taken with the acronym "TCP", which within the context of the rest of this document means "Termination Connection Point" and does not indicate the Transmission Control Protocol.

2.1. Terminology

The reader is assumed to be familiar with the terminology in [LMP] and [LMP-TEST]. The following ITU-T terminology/abbreviations are used in this document:

Connection Point (CP): A "reference point" that consists of a pair of co-located "unidirectional connection points" and therefore represents the binding of two paired bidirectional "connections".

Connection Termination Point (CTP): A connection termination point represents the state of a CP [M.3100].

Characteristic Information: Signal with a specific format, which is transferred on "network connections". The specific formats will be defined in the technology-specific recommendations. For trails, the Characteristic Information is the payload plus the overhead. The information transferred is characteristic of the layer network.

Link: A subset of ports at the edge of a subnetwork or access group that are associated with a corresponding subset of ports at the edge of another subnetwork or access group.

Link Connection (LC): A transport entity that transfers information between ports across a link.

Network Connection (NC): A concatenation of link and subnetwork connections.

Subnetwork: A set of ports that are available for the purpose of routing 'characteristic information'.

Subnetwork Connection (SNC): A flexible connection that is set up and released using management or control plane procedures.

Subnetwork Point (SNP): SNP is an abstraction that represents an actual or potential underlying connection point (CP) or termination connection point (TCP) for the purpose of control plane representation.

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

Termination Connection Point (TCP): A reference point that represents the output of a Trail Termination source function or the input to a Trail Termination sink function. A network connection represents a transport entity between TCPs.

Trail Termination source/sink function: A "transport processing function" that accepts the characteristic information of the layer network at its input, removes the information related to "trail" monitoring, and presents the remaining information at its output.

Unidirectional Connection: A "transport entity" that transfers information transparently from input to output.

Unidirectional Connection Point: A "reference point" that represents the binding of the output of a "unidirectional connection" to the input of another "unidirectional connection".

2.2. Abbreviations

LMP: Link Management Protocol

OTN: Optical Transport Network

PDH: Plesiosynchronous Digital Hierarchy

SDH: Synchronous Digital Hierarchy

SONET: Synchronous Optical Network

3. Transport Network Architecture

A generic functional architecture for transport networks is defined in International Telecommunication Union (ITU-T) Recommendation [G.805]. This recommendation describes the functional architecture of transport networks in a technology-independent way. This architecture forms the basis for a set of technology-specific architectural recommendations for transport networks (e.g., SDH, PDH, OTN, etc.).

The architecture defined in G.805 is designed using a layered model with a client-server relationship between layers. The architecture is recursive in nature; a network layer is both a server to the client layer above it and a client to the server layer below it. There are two basic building blocks defined in G.805: "subnetworks" and "links". A subnetwork is defined as a set of ports that are available for the purpose of routing "characteristic information". A link consists of a subset of ports at the edge of one subnetwork (or "access group") and is associated with a corresponding subset of ports at the edge of another subnetwork or access group.

Two types of connections are defined in G.805: link connection (LC) and subnetwork connection (SNC). A link connection is a fixed and inflexible connection, while a subnetwork connection is flexible and is set up and released using management or control plane procedures. A network connection is defined as a concatenation of subnetwork and link connections. Figure 1 illustrates link and subnetwork connections.




Figure 1: Subnetwork and Link Connections

G.805 defines a set of reference points for the purpose of identification in both the management and the control planes. These identifiers are NOT required to be the same. A link connection or a subnetwork connection is delimited by connection points (CPs). A network connection is delimited by a termination connection point (TCP). A link connection in the client layer is represented by a pair of adaptation functions and a trail in the server layer network. A trail represents the transfer of monitored adapted characteristics information of the client layer network between access points (APs).

A trail is delimited by two access points, one at each end of the trail. Figure 2 shows a network connection and its relationship with link and subnetwork connections. Figure 2 also shows the CP and TCP reference points.




Figure 2: Network Connection with Link and Subnetwork Connections

For management plane purposes, the G.805 reference points are represented by a set of management objects described in ITU-T Recommendation M.3100 [M.3100]. Connection termination points (CTPs) and trail termination points (TTPs) are the management plane objects for CP and TCP, respectively.

In the same way as in M.3100, the transport resources in G.805 are identified for the purposes of the control plane by entities suitable for connection control. G.8080 introduces the reference architecture for the control plane of the Automatically Switched Optical Networks (ASONS). G.8080 introduces a set of reference points relevant to the ASON control plane and their relationship to the corresponding points in the transport plane. A subnetwork point (SNP) is an abstraction that represents an actual or potential underlying CP or an actual or potential TCP. A set of SNPs that are grouped together for the purpose of routing is called SNP pool (SNPP). Similar to LC and SNC, the SNP-SNP relationship may be static and inflexible (this is referred to as an SNP link connection), or it can be dynamic and flexible (this is referred to as an SNP subnetwork connection).

3.1. G.8080 Discovery Framework

G.8080 provides a reference control plane architecture based on the descriptive use of functional components representing abstract entities and abstract component interfaces. The description is generic, and no particular physical partitioning of functions is implied. The input/output information flows associated with the functional components serve for defining the functions of the components and are considered to be conceptual, not physical. Components can be combined in different ways, and the description is not intended to limit implementations. Control plane discovery is described in G.8080 by using three components: Discovery Agent (DA), Termination and Adaptation Performer (TAP), and Link Resource Manager (LRM).

The objective of the discovery framework in G.8080 is to establish the relationship between CP-CP link connections (transport plane) and SNP-SNP link connections (control plane). The fundamental characteristics of G.8080 discovery framework is the functional separation between the control and the transport plane discovery processes and name spaces. From G.8080: "This separation allows control plane names to be completely separate from transport plane names, and completely independent of the method used to populate the DAs with those transport names. In order to assign an SNP-SNP link connection to an SNPP link, it is only necessary for the transport name for the link connection to exist". Thus, it is possible to assign link connections to the control plane without the link connection being physically connected.

Discovery encompasses two separate processes: (1) transport plane discovery, i.e., CP-to-CP and TCP-to-TCP connectivity; and (2) control plane discovery, i.e., SNP-to-SNP and SNPP links.

G.8080 Amendment 1 defines the Discovery Agent (DA) as the entity responsible for discovery in the transport plane. The DA operates in the transport name space only and in cooperation with the Termination and Adaptation Performer (TAP), provides the separation between that space and the control plane names. A local DA is only aware of the CPs and TCPs that are assigned to it. The DA holds the CP-CP link connection in the transport plane to enable SNP-SNP link connections to be bound to them at a later time by the TAP. The CP-CP relationship may be discovered (e.g., per G.7714.1) or provided by a management system.

Control plane discovery takes place entirely within the control plane name space (SNPs). The Link Resource Manager (LRM) holds the SNP-SNP binding information necessary for the control plane name of the link connection, while the termination adaptation performer (TAP) holds

the relation between the control plane name (SNP) and the transport plane name (CP) of the resource. Figure 3 shows the relationship and the different entities for transport and control discoveries.

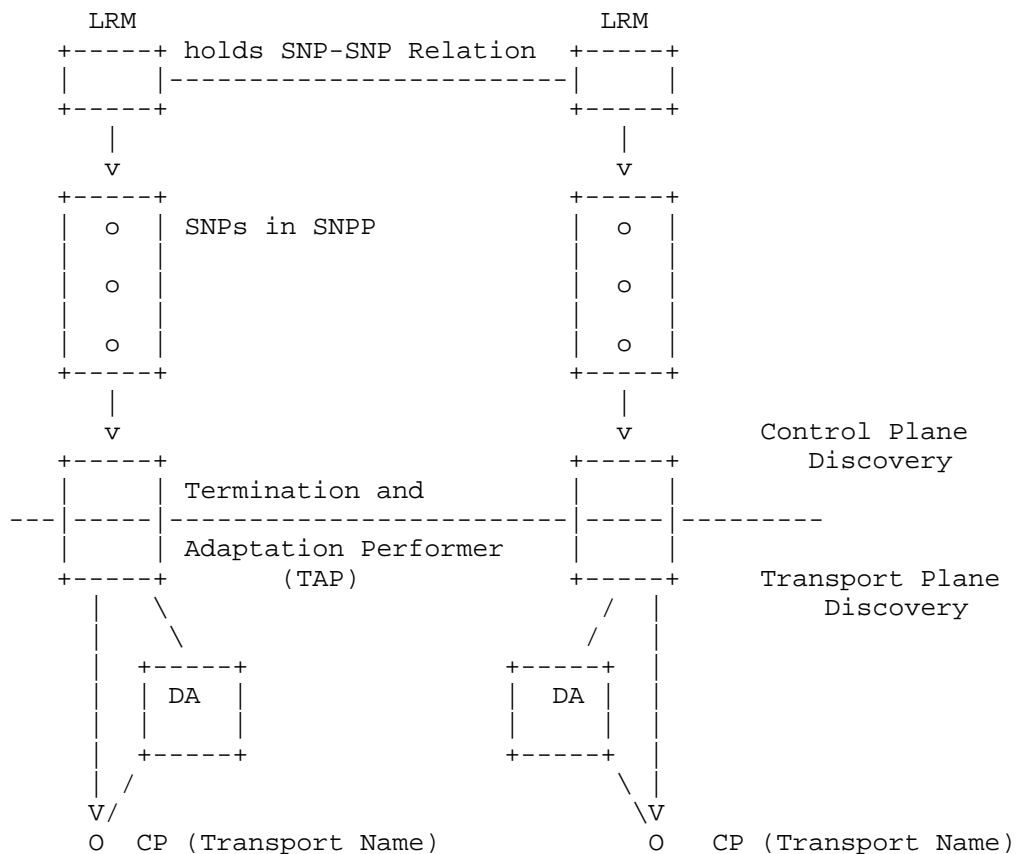


Figure 3: Discovery in the Control and the Transport Planes

4. Discovery Technologies

4.1. Generalized Automatic Discovery Techniques G.7714

Generalized automatic discovery techniques are described in G.7714 to aid resource management and routing for G.8080. The term routing here is described in the transport context of routing connections in an optical network as opposed to the routing context typically associated in packet networks.

G.7714 is concerned with two types of discovery:

- Layer adjacency discovery
- Physical media adjacency discovery

Layer adjacency discovery can be used to correlate physical connections with management configured attributes. Among other features this capability allows reduction in configuration and the detection of mis-wired equipment.

Physical media adjacency discovery is a process that allows the physical testing of the media for the purpose of inventory capacity and verifying the port characteristics of physical media adjacent networks.

G.7714 does not specify specific protocols but rather the type of techniques that can be used. G.7714.1 specifies a protocol for layer adjacency with respect to SDH and OTN networks for layer adjacency discovery. A GMPLS method for layer discovery using elements of LMP is included in this set of procedures.

An important point about the G.7714 specification is that it specifies a discovery mechanism for optical networks but not necessarily how the information will be used. It is intended that the transport management plane or a transport control plane may subsequently make use of the discovered information.

4.2. LMP and G.8080 Terminology Mapping

GMPLS is a set of IP-based protocols, including LMP, providing a control plane for multiple data plane technologies, including optical/transport networks and their resources (i.e., wavelengths, timeslots, etc.) and without assuming any restriction on the control plane architecture (see [RFC3945]). On the other hand, G.8080 defines a control plane reference architecture for optical/transport networks without any restriction on the control plane implementation. Being developed in separate standards forums, and with different scopes, they use different terms and definitions.

Terminology mapping between LMP and ASON (G.805/G.8080) is an important step towards the understanding of the two architectures and allows for potential cooperation in areas where cooperation is possible. To facilitate this mapping, we differentiate between the two types of data links in LMP. According to LMP, a data link may be considered by each node that it terminates on as either a 'port' or a 'component link'. The LMP notions of port and component link are supported by the G.805/G.8080 architecture. G.8080's variable adaptation function is broadly equivalent to LMP's component link, i.e., a single server-layer trail dynamically supporting different multiplexing structures. Note that when the data plane delivers its own addressing space, LMP Interface_IDs and Data Links IDs are used as handles by the control plane to the actual CP Name and CP-to-CP Name, respectively.

The terminology mapping is summarized in the following table: Note that the table maps ASON terms to GMPLS terms that refer to equivalent objects, but in many cases there is not a one-to-one mapping. Additional information beyond discovery terminology can be found in [LEXICO].

ASON Terms	GMPLS/LMP Terms Port	GMPLS/LMP Terms Component Link
CP	TE Resource; Interface (Port)	TE Resource; Interface. (Comp. link)
CP Name	Interface ID no further sub- division for(label) resource allocation	Interface ID(s) resources (such as timeslots, etc.) on this interface are identified by set of labels
CP-to-CP Link	Data Link	Data Link
CP-to-CP Name	Data Link ID	Data Link ID
SNP	TE Resource	TE Resource
SNP Name	Link ID	Link ID
SNP LC	TE Link	TE Link
SNP LC Name	TE Link ID	TE Link ID
SNPP	TE Link End (Port)	TE Link End (Comp. Link)
SNPP Name	Link ID	Link ID
SNPP Link	TE Link	TE Link
SNPP Link Name	TE Link ID	TE Link ID

where composite identifiers are:

- Data Link ID: <Local Interface ID; Remote Interface ID>
- TE Link ID: <Local Link ID; Remote Link ID>

Composite Identifiers are defined in the RFC 4204 [LMP]. LMP discovers data links and identifies them by the pair of local and remote interface IDs. TE links are composed of data links or component TE links. TE links are similarly identified by pair of local and remote link ID.

4.2.1. TE Link Definition and Scope

In the table, TE link/resource is equated with the concept of SNP, SNP LC, SNPP, and SNPP link. The definition of the TE link is broad in scope, and it is useful to repeat it here. The original definition appears in [GMPLS-RTG]:

"A TE link is a logical construct that represents a way to group/map the information about certain physical resources (and their properties) that interconnects LSRs into the information that is used by Constrained SPF for GMPLS path computation, and GMPLS signaling".

While this definition is concise, it is probably worth pointing out some of the implications of the definition.

A component of the TE link may follow different paths between the pair of LSRs. For example, a TE link comprising multiple STS-3cs, the individual STS-3cs component links may take identical or different physical (OC-3 and/or OC-48) paths between LSRs.

The TE link construct is a logical construction encompassing many layers in networks [RFC3471]. A TE link can represent either unallocated potential or allocated actual resources. Further allocation is represented by bandwidth reservation, and the resources may be real or, in the case of packets, virtual to allow for overbooking or other forms of statistical multiplexing schemes.

Since TE links may represent large numbers of parallel resources, they can be bundled for efficient summarization of resource capacity. Typically, bundling represents a logical TE link resource at a particular Interface Switching Capability. Once TE link resources are allocated, the actual capacity may be represented as LSP hierarchical (tunneled) TE link capability in another logical TE link [HIER].

TE links also incorporate the notion of a Forwarding Adjacency (FA) and Interface Switching Capability [RFC3945]. The FA allows transport resources to be represented as TE links. The Interface Switching Capability specifies the type of transport capability such as Packet Switch Capable (PSC), Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber-Switch Capable (FSC).

A TE link between GMPLS-controlled optical nodes may consist of a bundled TE link, which itself consists of a mix of point-to-point component links [BUNDLE]. A TE link is identified by the tuple (link Identifier (32-bit number), Component link Identifier (32-bit number), and generalized label (media specific)).

4.3. LMP and G.8080 Discovery Relationship

LMP currently consists of four primary procedures, of which the first two are mandatory and the last two are optional:

1. Control channel management
2. Link property correlation
3. Link verification
4. Fault management

LMP procedures that are relevant to G.8080 control plane discovery are control channel management, link property correlation, and link verification. Key to understanding G.8080 discovery aspects in relation to [LMP] is that LMP procedures are specific for an IP-based control plane abstraction of the transport plane.

LMP control channel management is used to establish and maintain control channel connectivity between LMP adjacent nodes. In GMPLS, the control channels between two adjacent nodes are not required to use the same physical medium as the TE links between those nodes. The control channels that are used to exchange the GMPLS control plane information exist independently of the TE links they manage (i.e., control channels may be in-band or out-of-band, provided the associated control points terminate the LMP packets). The Link Management Protocol [LMP] was designed to manage TE links, independently of the physical medium capabilities of the data links.

Link property correlation is used to aggregate multiple data links into a single TE link and to synchronize the link properties.

Link verification is used to verify the physical connectivity of the data links and verify the mapping of the Interface-ID to Link-ID (CP to SNP). The local-to-remote associations can be obtained using a priori knowledge or using the link verification procedure.

Fault management is primarily used to suppress alarms and to localize failures. It is an optional LMP procedure; its use will depend on the specific technology's capabilities.

[LMP] supports distinct transport and control plane name spaces with the (out-of-band) TRACE object (see [LMP-TEST]). The LMP TRACE object allows transport plane names to be associated with interface identifiers [LMP-TEST].

Aspects of LMP link verification appear similar to G.7714.1 discovery; however, the two procedures are different. G.7714.1 provides discovery of the transport plane layer adjacencies. It provides a generic procedure to discover the connectivity of two

endpoints in the transport plane. On the other hand, the LMP link verification procedure is a control-plane-driven procedure and assumes either (1) a priori knowledge of the associated data plane's local and remote endpoint connectivity and Interface_IDs (e.g., via management plane or use of G.7714.1), or (2) support of the remote node for associating the data interface being verified with the content of the TRACE object (inferred mapping). For SONET/SDH transport networks, LMP verification uses the SONET/SDH Trail Trace identifier (see [G.783]).

G.7714.1 supports the use of transport plane discovery independent of the platform using the capability. Furthermore, G.7714.1 specifies the use of a Discovery Agent that could be located in an external system and the need to support the use of text-oriented man-machine language to provide the interface. Therefore, G.7714.1 limits the discovery messages to printable characters defined by [T.50] and requires Base64 encoding for the TCP-ID and DA ID. External name-servers may be used to resolve the G.7714.1 TCP name, allowing the TCP to have an IP, Network Service Access Protocol (NSAP), or any other address format. On the other hand, LMP is based on the use of an IP-based control plane, and the LMP interface ID uses IPv4, IPv6, or unnumbered interface IDs.

4.4. Comparing LMP and G.8080

LMP exists to support GMPLS TE resource and TE link discovery. In section 4.2.1, we elaborated on the definition of the TE link. LMP enables the aspects of TE links to be discovered and reported to the control plane, more specifically, the routing plane. G.8080 and G.7714 are agnostic to the type of control plane and discovery protocol used. LMP is a valid realization of a control plane discovery process under a G.8080 model.

G.7714 specifies transport plane discovery with respect to the transport layer CTPs or TCPs using ASON conventions and naming for the elements of the ASON control plane and the ASON management plane. This discovery supports a centralized management model of configuration as well as a distributed control plane model; in other words, discovered items can be reported to the management plane or the control plane. G.7714.1 provides one realization of a transport plane discovery process.

Today, LMP and G.7714, G7714.1 are defined in different standards organizations. They have evolved out of different naming schemes and architectural concepts. Whereas G.7714.1 supports a transport plane layer adjacency connectivity verification that can be used by a

control plane or a management plane, LMP is a control plane procedure for managing GMPLS TE links (GMPLS's control plane representation of the transport plane connections).

5. Security Considerations

Since this document is purely descriptive in nature, it does not introduce any security issues.

G.8080 and G.7714/G.7714.1 provide security as associated with the Data Communications Network on which they are implemented.

LMP is specified using IP, which provides security mechanisms associated with the IP network on which it is implemented.

6. Informative References

- [LMP] Lang, J., "Link Management Protocol (LMP)", RFC 4204, October 2005.
- [LMP-TEST] Lang, J. and D. Papadimitriou, "Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) Encoding for Link Management Protocol (LMP) Test Messages", RFC 4207, October 2005.
- [RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [GMPLS-RTG] Kompella, K. and Y. Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [HIER] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.
- [BUNDLE] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.

- [LEXICO] Bryskin, I. and A. Farrel, "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", Work in Progress, January 2006.

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

- [G.783] ITU-T G.783 (2004), Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks.
- [G.805] ITU-T G.805 (2000), Generic functional architecture of transport networks.
- [G.7714] ITU-T G.7714/Y.1705 (2001), Generalized automatic discovery techniques.
- [G.7714.1] ITU-T G.7714.1/Y.1705.1 (2003), Protocol for automatic discovery in SDH and OTN networks.
- [G.8080] ITU-T G.8080/Y.1304 (2001), Architecture for the automatically switched optical network (ASON).
- [M.3100] ITU-T M.3100 (1995), Generic Network Information Model.
- [T.50] ITU-T T.50 (1992), International Reference Alphabet.

7. Acknowledgements

The authors would like to thank Astrid Lozano, John Drake, Adrian Farrel and Stephen Shew for their valuable comments.

The authors would like to thank ITU-T Study Group 15 Question 14 for their careful review and comments.

Authors' Addresses

Don Fedyk
Nortel Networks
600 Technology Park Drive
Billerica, MA, 01821

Phone: +1 978 288-3041
EMail: dwfedyk@nortel.com

Osama Aboul-Magd
Nortel Networks
P.O. Box 3511, Station 'C'
Ottawa, Ontario, Canada
K1Y-4H7

Phone: +1 613 763-5827
EMail: osama@nortel.com

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

EMail: dbrungard@att.com

Jonathan P. Lang
Sonos, Inc.
223 E. De La Guerra
Santa Barbara, CA 93101

EMail: jplang@ieee.org

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

Phone: +32 3 240-84-91
EMail: dimitri.papadimitriou@alcatel.be

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).