Status and Trends for Standardization of Architecture and Solutions for Multi-Domain Optical Networks.

Adrian Farrel[1]
[1]Old Dog Consulting

Abstract

Optical technologies form the key foundations for transport networks used by all major telecommunication carriers. Those networks are made up of equipment supplied by different manufacturers, and utilize different optical switching and transmission technologies. These give rise to distinct islands or "domains" of network nodes of different capabilities. Further, administrative segmentation of networks, and inter-carrier communications add another concept of a domain.

Control plane technologies have seen increasing popularity as a way of discovering transport network capabilities and provisioning connectivity across them. To date, the focus has been on control plane operations within optical domains, but attention is now turning to the operation of multi-domain optical networks.

This paper sets out the current status and trends in standardization of control plane architectures and technologies for use in multi-domain optical networks.

Key Words - Multi-domain, Multi-layer, Optical Networking, GMPLS, ASON, Standardization, IETF, ITU-T.

I. Introduction

The introduction of fiber optics in the early 1980s heralded a revolution in transport networks. Over time, physical technologies have been standardized in support of Time Division Multiplexing (TDM) and Wave Division Multiplexing (WDM) to make it possible to use fiber optical to interconnect equipment from different vendors to build optical networks. Significant networks are now built from optical equipment with hundred of network elements and fiber spans of thousands of kilometers.

The management and operation of optical networks used to be a significant undertaking. An optical trail (a traffic path through the network) was created by configuring each optical node along the path to enable the interfaces and enable cross-connects. Such configuration was achieved using Element Management Systems (EMSs) under the control of a centralized Network Management System (NMS) or Operations Service System (OSS).

The development of dynamic control plane technologies offered considerable simplifications to the operation of networks and has a long history in computer networking. In the late 1990s, the Internet Engineering Task Force (IETF) started work on a series of extensions to the Multiprotocol Label Switching (MPLS) protocol suite to produce Generalized MPLS (GMPLS) [1], a set of control protocols suitable for establishing trails through optical networks. At roughly the same time, the International Telecommunications Union (ITU) began to formally document the architecture of transport networks and to create a specific architecture for the Automatically Switched Optical Network (ASON) [2].

The benefits afforded by the use of a dynamic control plane and the potential that it offers in terms of dynamic network operation, recovery from faults, reduced operational complexity, rapid service provisioning, ability to realize revenue from existing equipment, and mechanisms for traffic engineering are well rehearsed and will not be repeated in this paper. Suffice to say that control planes for optical networks have attracted sufficient attention that they are now routinely a consideration in the procurement and deployment of optical equipment.

As optical networks grow in size, it becomes natural to partition them. The motivations may be the same as for any other network (administrative borders, network mergers through acquisitions, or scaling concerns), but the optical technologies also give rise to other reasons. For example, it may be possible to provide connectivity for one optical network (say a TDM network) by provisioning trails in another network (perhaps a WDM network). This network layering is fundamental to the ASON architecture [1] and is also recognized within the IETF [3]. In both cases, the border between the layers serves to partition the network into separate domains and creates a client/server relationship between the domains.

A further motivation for segmenting the network may arise from the optical equipment itself. Different vendors may choose to enhance their offering with special features and advanced functions that are not part of the standardized data plane or control plane. In order to construct a larger network, an operator must arrange vendor equipment into islands and interconnect those islands to the rest of the network.

This paper examines the current state of affairs in the relevant standards bodies with respect to architectures and solutions for multi-domain optical networks. It introduces standard techniques for provisioning and routing in optical networks, and then discusses and defines the optical network domain. The subsequent sections describe the work in the standard bodies for the distinct issues raised by inter-layer and inter-domain networking.

II. Standardized Provisioning and Routing in Optical Networks

The optical control plane consists of protocols for discovering optical connectivity, distributing information about available network resources, and provisioning optical trails.

The IETF's Link Management Protocol (LMP) [4] provides a network node with mechanisms to establish the connectivity and nature of the data links to and from its neighbor. The data links can be verified using LMP's test procedures, and faults can be isolated through control plane exchanges – a feature that is particularly useful in transparent optical networks where there is no in-fiber Operations and Management (OAM) function. LMP also allows data links to be clustered into administrative units called Traffic Engineering links (TE links).

LMP's capabilities are enhanced by the ITU's link discovery procedures defined in G.7714.1 [5] that enable 'plug and play' function in optical networks so that it is not necessary to configure the identity of neighbors when new fibers or data links are connected to a node. These protocol procedures are based on the IETF's Point-to-Point Protocol Link Control Protocol extensions [6].

Routing protocols in the optical control plane are used to distribute link state information about each TE link. Each node advertises the capabilities of the links it terminates and states the available resources (bandwidth) on the links. As network resources are used or released, the node re-advertises the links with updated resource availability. Thus, every node in the network is able to construct a Traffic Engineering Database (TED) providing a full and up-to-date view of the of topology of the network and the available potential paths through the network. The IETF has specified two routing protocols for use in optical networks: GMPLS-OSPF [7] provides protocol extensions to the well-known IP routing protocol, OSPF; GMPLS-IS-IS [8] makes similar extensions to the IS-IS protocol.

Optical trails are called Label Switched Paths (LSPs) in GMPLS; the labels in this context are physical resources in the optical medium such as timeslots or wavelengths. LSPs are across the network established by signaling protocols. Initially, the IETF defined two functionally equivalent protocols for use in optical networks: Constraint-based Routed Label Distribution Protocol (CR-LDP) [7], and the Resource Reservation Protocol TE Extensions (RSVP-TE) [8]. Both of these use common protocol elements [9], however, recognizing that the industry does not benefit from the existence of two equivalent protocols, the IETF abandoned CR_LDP and progressed only with RSVP-TE [10].

Routing and signaling specifications in the ITU have so far focused on the requirements at domain boundaries. This topic is discussed further in section V.

Standardization of optical control plane architecture and protocols continues in both the IETF and the ITU. The ITU, with a stronger participation from operators has a heavier bias towards architecture and deployment models, while the IETF is more focused on protocol design, implementation and 'running code.'

III. Introducing Multi-Layer and Multi-Domain Networks

Current optical networks span vast distances and encompass many network nodes. Nodes that share a common management policy or addressing scheme are considered to form domains. Domains may also comprise network elements of the same switching type (that is, transport technology). Administrative subdivisions create domains for commercial reasons or for to achieve scalability and management simplifications within a network. A server network that provides connectivity for one or more client networks represents a separate domain. Clusters or islands of nodes with specific or proprietary control plane behavior (including subnetworks that don't use a control plane at all) also form domains.

In all of these cases, a domain may be classified as "any collection of network elements within a common sphere of address management or path computational responsibility" [11]. This definition fits well with the concepts of routing areas and Autonomous Systems (ASes) familiar in Internet routing, and also matches the idea of an ASON Routing Area (RA) [2] and [12], and is coherent with the subnetwork defined in ASON [2].

Domains of nodes of the same switching technology form special domain types known in the IETF as regions [3]. IETF regions may comprise WDM nodes, TDM nodes, Layer 2 switching nodes, or packet switching nodes: in the context of this paper, we are only interested in WDM and TDM switching. Regions may be further subdivided into layers or sub-layers according to the capabilities of the switching type. For example, the TDM region may be subdivided into VC4 and VC3 layers. Multi-layer networks are recognized by both the IETF [3] and the ITU [2].

Multi-domain and multi-layer networking present very different, yet in many ways similar architectural challenges and require consistent protocol solutions.

IV. Architectures for Multi-Domain and Multi-Layer Networks

A well-defined architectural model is essential to the correct development of protocols and their proper deployment. The ITU has included multiple domains in their optical network architecture from the start while the IETF's approach ahs been less formal and driven by specific implementation and deployment needs. Thus the documentation of the IETF architecture for optical networks lags behind that of the ITU, but protocol solutions from the two bodies can be successfully combined to produce a high-function solution. To do this, it is important to understand the basics of the architectural models, and these are introduced below.

Peer Domains

The most basic domains operate as peers. In the Internet world, we may consider ASes as the prime example of peer domains: they operate on an equal footing, exchange no TE routing information, yet must cooperate to establish end-to-end LSPs. Requirements for inter-AS traffic engineering are described in RCF 4216 [13].

The basic unit of architecture in the ASON model is the subnetwork. Thus, a single network node may be considered as a subnetwork; a self-contained domain in its own right. But this may be no more than an academic distinction; it is more interesting to examine the subnetwork as a collection of nodes that itself has connections to the outside world. ASON subnetworks may operate as peer domains, so a network may be constructed as shown in Figure 1.

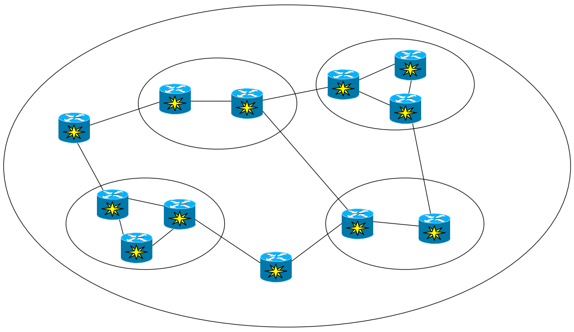

Figure 1 : The basic multi-domain architecture

It can be seen that the ASON multi-domain network is more open than a network built of ASes. But this should not be a surprise, they are intended to operate on a wholly different scale with ASON subnetworks potentially containing just a few network elements while ASes are usually large collections of very many nodes.

Hierarchical Domains

As shown in Figure 1, a network may be constructed from a collection of subnetworks. The ASON architecture makes this model recursive. That is, a subnetwork may itself be constructed from a collection of subnetworks giving a fully-featured abstract architectural model as shown in Figure 2.

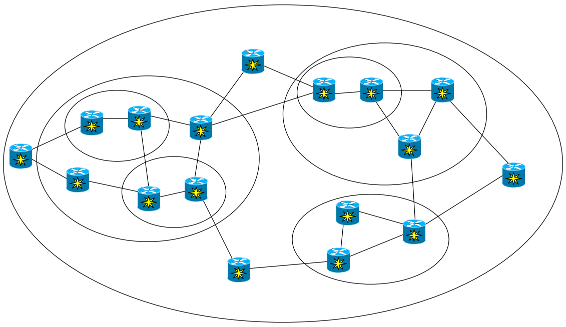

Figure 2 : The full ASON architectural model

Routing Areas and Routing Levels

Since the objective is to establish end-to-end connectivity, and since the control plane needs TE routing information to achieve this, we need to determine how the architectures accommodate routing.

The ASON architecture defines the concept of Routing Areas (RAs). According to G.8080 an RA is defined by a set of subnetworks, the links that interconnect them, and the ends of the links that connect to the outside world [2]. This makes an RA look very much like a subnetwork and, indeed RAs can be arbitrarily nested with the limit of subdivision being an RA containing just two subnetworks and one link (where a subnetwork may itself be just a single node). These hierarchical arrangements of RAs are termed Routing Levels.

The Internet has a similar concept in IP routing. Routing areas in OSPF and levels in IS-IS, may be arranged hierarchically, although by convention in IS-IS and by definition in OSPF this relationship is never stacked more than two deep.

Multi-Layer Networks

Multi-layer networks also fit the multi-domain model. In an ASON network it is as simple as recognizing that a layer is contained within its parent layer and may be treated as a subnetwork or routing area. Indeed, G.8080 states that an RA exists within a single layer network indicating the implicit containment relationship as shown in Figure 3.

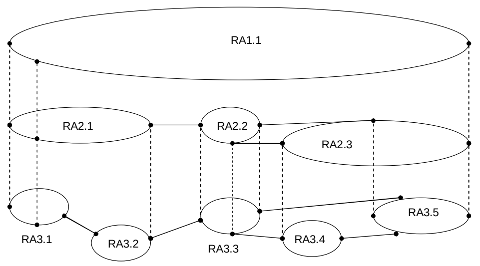

Figure 3 : The ASON architecture for routing levels showing hw routing areas and external points of connectivity are mapped into the routing areas of higher levels.

GMPLS takes a slightly different view of the multi-layer network. Nodes of different technologies may be present in a single routing area under the control of a single instance of a GMPLS Routing protocol as described in [3]. On the other hand, a more common GMPLS multi-layer architectural model presents the layer networks in a client/server relationship as shown in Figure 4. In this case, there is no routing information exchange between the network layers, but connections across the lower layer network may be presented as *virtual links* in the upper layer network.

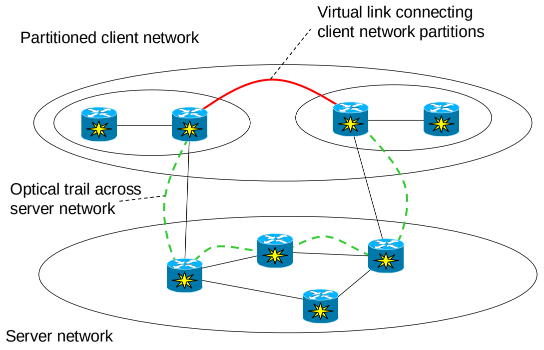

Figure 4 : The client/server multi-layer network.

Of course, the ASON architecture is quite capable of accommodating the client/server multi-layer network.

The Virtual Network Topology

The IETF have formalized the client/server multi-layer network relationship in terms of a Virtual Network Topology (VNT) [xx]. The LSPs across the lower layer network are established to meet the needs of the client network. These lower layer LSPs provide virtual links in the client network and so form a virtual network topology. The connectivity may be established on-demand as LSPs in the higher layer are attempted, or, more probably, under the control and supervision of management applications. The latter facilitates a clear separation between the administration of the networks while keeping the management function very open to different interpretations of policy in different deployments.

ASON may also support a similar concept. Server layer trails me be provisioned to provide connectivity in the client layer, and this, it could be argued, is fundamental to the ASON multi-layer network.

V. Inter-Domain Interactions

So far, this paper has concentrated on the architectures of multi-domain networks. But these are abstract concepts useful for theory or network planning. What is also required is control plane protocols to establish end-to-end connectivity, and to exchange the routing information needed to determine paths across multi-domain optical networks.

Signaling

The signaling protocols are responsible for installing the optical trails, causing optical resources to be enabled and cross-connects (electronic or optical) to be programmed. An end-to-end trail can be constructed in one of three ways [11].

- *Contiguous* LSPs are formed of a single, end-to-end protocol exchange resulting in a coherent protocol ‘session’ from source to destination. This mode of operation is suitable for the case where the source or ingress node wishes to maintain end-to-end control of the LSP.
- LSP segments may be *stitched* [14] together to form end-to-end data plane connectivity, but with each segment under the control of the domain it crosses. This model is particularly suited to situations where each domain requires greater control of the connectivity that it provides and where the domains may use different mechanisms to deliver the end-to-end level of service.
- The *hierarchical LSP* [15] is used to ‘tunnel’ a client layer LSP over a server layer network and is, therefore, most applicable to the multi-layer form of the multi-domain network.

The ASON architecture defines three key interfaces, or *reference points* in the construction of multi-domain networks. The User Network Interface (UNI) exists between a user (such as a client network) and a network. The Internal Network-to-Network Interface (I-NNI) exists between network nodes in the same subnetwork, and the External Network-to-Network Interface (E-NNI) provides a for interaction between subnetworks. In some cases, the interface between client and server networks may be considered to be an E-NNI rather than a UNI, thus providing somewhat greater integration between the layer networks.

The ITU has developed three signaling specifications that offer the same level of abstract functionality for use both at the UNI and the E-NNI. These specifications build on existing signaling protocols and are: G.7713.1 utilizing the PNNI protocol developed for ATM; G.7713.2 adding extensions to the GMPLS RSVP-TE protocol; and G.7713.3 enhancing GMPLS CR-LDP for the same purpose. For various reasons, including the support of the Optical Interworking Forum (OIF), G,7713.2 [16] is the protocol that has seen most interest.

The IETF, too, recognizes the importance of the abstract reference points and has looked to see how its standard GMPLS signaling protocol [8] can be applied to the UNI. This resulted in RFC 4208 [17] that describes the applicability of GMPLS signaling to the UNI, and (although not recognized in the text of that document) also provides a description of how GMPLS RSVP-TE can operate at the E-NNI.

The ITU has not devoted any effort to signaling specifications for the I-NNI. Since one motivation for the existence of a subnetwork is to allow proprietary implementations, this makes some sense. But the IETF, on the other hand, is concerned to achieve full control plane interoperability between devices and so the primary focus of GMPLS has been the I-NNI.

Routing

Routing exchanges between domains have provided a rich seam of discussion. If domains exist in order to facilitate administrative boundaries, why would they exchange routing information? If the purpose of domains is to make the network more scalable, doesn’t routing exchange defeat the objective? And if domains are built from distinct proprietary implementations, what routing information can they meaningfully exchange?

Yet in order to achieve an end-to-end connection, routing decisions must be made. The signaling protocol must be instructed what path to take, and these instructions must be based on some knowledge of the connectivity of the network and the available network resources.

In the client/server model, the ASON UNI very specifically forbids the exchange of routing information [2]. This means that the client network cannot know whether the server network is able to provide connectivity – it must simply make a request and wait to hear whether the request is successful. That limitation means that the VNT model with pre-provisioned server LSPs presented as virtual links in the client network is most applicable. Attempts to represent potential connectivity across server networks into the client network as a mesh of virtual links or by showing the server network as a virtual node lead, inevitably, to the many inaccuracies and scaling complications of TE aggregation.

The E-NNI, on the other, hand allows controlled ‘leakage’ of routing information from one domain to another. For this reason, the E-NNI is often considered to be applicable between client/server networks as well as between peer networks. In fact, of course, what we are interested in here is the exchange of information between Routing Areas, and since the relationship between RAs is strictly hierarchical [12] the information exchange is also hierarchical.

Instead of developing its own protocol solution for inter-level routing exchange, the ITU has agreed to work with the IETF to develop a suitable set of extensions to GMPLS OSPF. Working from requirements [18] and an analysis of the existing capabilities of OSPF, a design team made of ITU and IETF participants have evolved OSPF capabilities capable of providing the necessary information exchange [19]. However, exactly what information is exchanged and under what circumstances remains an open issue with significant concerns about scalability and TE aggregation being expressed within the IETF.

Path Computation Element

An alternative approach to routing in multi-domain networks has been developed in the IETF. Instead of relying on information distribution to allow a source node to compute an end-to-end path, this mechanism distributes the path computation request.

In the Path Computation Element (PCE) architecture [20] a Path Computation Client (PCC) such as the head-end node of an LSP makes a request to a PCE in its own domain for an end-to-end path. The PCE may consult with other PCEs in other domains to determine the most suitable route through the network.

The PCE architecture is highly applicable to multi-area and multi-AS environments [20] and several different modes of cooperation between signaling and PCE have been defined to handle various concerns including scaling, domain selection, and domain confidentiality.

Work is also underway to examine how PCE may be applied to multi-layer networks [21] both in terms of simple end-to-end path derivation, and in relation to the VNT construct.

At the same time, the ITU have embraced PCE as a possible solution to the ASON routing architecture. Recommendation G.7715.2 [22] shows how PCE may be used to provide *routing controller* function in the Routing Areas of and ASON network.

Calls

This paper has concentrated on a description of how end-to-end data plane connectivity is provided. But the ITU's ASON architecture also includes the concept of a *call* that provides coordination of service provision between reference points within the network.

Figure 5 shows how a calls relate to connections in the network. It can be seen that in order to provide the end-to-end connection, a series of connection segments are stitched together. In some subnetworks multiple, parallel segments are needed to achieve the level of function that can be provided in other subnetworks by a single segment. Coordinating the whole process and helping to stitch the segments together is the call, itself made up of segments. The call also has the more traditional purpose of admitting the service both at the destination and at the transit E-NNI reference points.

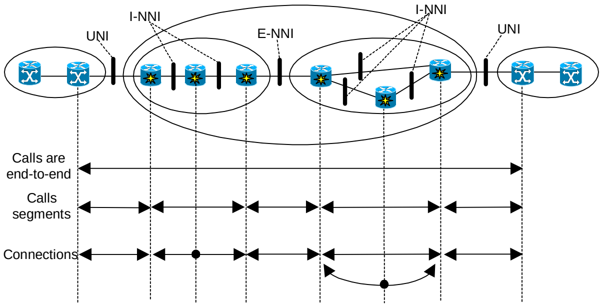

Figure 5 : Calls and connections in the ASON

The ITU's signaling protocol extensions previously described provide support for calls. They allow a UNI client to request a connection and associated call, and for this call to be presented at each E-NNI reference point on the path before final delivery at the remote UNI. In the signaling protocol presented in [16] calls are 'piggy-backed' on connection setup messages, and this, while functional for simple calls and connections, is not so easy to flexibly extend to scenarios with multiple connections associated with a single call, especially where those connections are diversely routed through the network making use of different subnetworks. Conversely, the inability to separate calls and connections in this way makes it impossible to set up a call without also defining an associated connection.

For many years, the IETF's GMPLS specifications completely ignored calls. They were an architectural construct that was not needed for early implementations and deployments and so they were not factored into any protocol specifications. It is only relatively recently, reacting to calls from the ITU for the IETF to develop GMPLS to meet the full set of ASON signaling requirements [23] that the Common Control and Measurement Plane (CCAMP) working group of the IETF specified how to achieve calls in GMPLS.
GMPLS calls use a significantly different protocol mechanism [24] from that used in ASON signaling, In part

this is a consequence of a need to satisfy the full set of ASON signaling requirements including the need for call/connection separation and for multiple connections associated with a single call.

In practice, the two call signaling mechanisms could interoperate successfully since the GMPLS mechanism can be used to carry the call across a subnetwork or domain, while the ASON technique can be applied at the UNI and E-NNI reference points.

VI. Conclusions and Ongoing Work

This paper has shown that there are many existing standards in place from the IETF and ITU-T for multi-domain optical networks. Both the multi-domain and multi-layer networks are covered by suitable architectural specifications, and protocol solutions exist to address many of the requirements for these network deployments.

Some deficiencies still exist in the advanced areas of function such as call/connection separation and call routing, but the main item of work outstanding for the standards bodies is a harmonization of signaling protocols. The industry does not benefit from multiple specifications in the same technical space, and it is incumbent on the leadership and technical membership of both the ITU and the IETF to develop methods for interoperability between the protocols that have been defined and to ensure convergence between the solutions through elimination of all differences.

As the control plane technologies for multi-domain optical mature and become more widely implemented and deployed, there will be a need for simplification. By removal of features that are determined to be unwanted or over-specified, the protocol standards and architectures will become more relevant and useful. More robust implementations will follow, and this will be of benefit to the whole industry.

Lastly, in an attempt to build substantial and useful multi-domain and multi-layer networks, further careful analysis of TE aggregation will be required in order to ensure that the networks are built on sound and scalable principles.

VII. References

[1] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004, IETF.
[2] "Architecture for the automatically switched optical network (ASON)", G.8080, November 2001, ITU-T.
[3] Shiomoto, K., Papadimitriou, D., Le Roux, J.L., Vigoureux, M., Brungard, D., "Requirements for GMPLS-based multi-region and multi-layer networks", draft-ietf-ccamp-gmpls-mrn-reqs-07, November 2007, work in progress, IETF.
[4] Lang, J., "Link Management Protocol (LMP)", RFC 4204, October 2005, IETF.
[5] "Protocol for automatic discovery in SDH and OTN networks", G.7714.1, April 2003, ITU-T.
[6] Simpson, W., "PPP LCP Extensions", RFC 1570, January 1994, IETF.
[7] Ashwood-Smith, P. and Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling - Constraint-based Routed Label Distribution Protocol (CR-LDP) Extensions", RFC 3472, January 2003, IETF.
[8] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling - Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003, IETF.
[9] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003, IETF.
[10] Andersson, L., and Swallow, G., "The Multiprotocol Label Switching (MPLS) Working Group decision on MPLS signaling protocols", RFC 3468, February 2003, IETF.
[11] Farrel, A., Vasseur, JP., and Ayyangar, A., "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006, IETF.
[12] "Architecture and requirements for routing in the automatically switched optical networks", G.7715, June 2002, ITU-T.
[13] Zhang, R., and Vasseur, JP., "MPLS Inter-AS Traffic Engineering Requirements", RFC 4216, November 2005, IETF.
[14] Ayyangar, A., Kompella, K., Vasseur, JP., and Farrel, A., "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", draft-ietf-ccamp-lsp-stitching-06.txt, April 2007, work in progress, IETF.
[15] Kompella, K. and Rekhter, Y., "LSP Hierarchy with Generalized MPLS TE", RFC 4206, October 2005, IETF.
[16] "Distributed Call and Connection Management: Signalling mechanism using GMPLS RSVP-TE", G.7713.2, March 2003, ITU-T.
[17] Swallow, G., Drake, J., Ishimatsu, H., and Rekhter, Y., "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005, IETF.
[18] Brungard, D., "Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Routing for the

Automatically Switched Optical Network (ASON)ASON routing requirements”, RFC 4258, November 2005, IETF.

[19] Papadimitriou, D., “OSPFv2 Routing Protocols Extensions for ASON Routing”, draft-ietf-ccamp-gmpls-ason-routing-ospf-04.txt, December 2007, work in progress, IETF.

[20] Farrel, A., Vasseur, JP., and Ash, G., “A Path Computation Element (PCE)-Based Architecture”, RFC 4655, August 2006, IETF.

[21] Oki, E., Le Roux, J-L., and Farrel, A., “Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering”, draft-ietf-pce-inter-layer-frwk-05.txt, September 2007, work in progress, IETF.

[22] “ASON routing architecture and requirements for remote route query”, February 2007, G.7715.2, ITU-T.

[23] Papadimitriou, D., Drake, J., Ash, G., Farrel, A., and Ong, L., “Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON) ASON signaling requirements”, RFC 4139, July 2005, IETF.

[24] Papadimitriou, D., and Farrel, A., “Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in Support of Calls GMPLS calls”, RFC 4974, August 2007, IETF.