Policy-Based Service Provisioning Architecture for Hybrid Photonic Networks

Belkacem Daheb, Guy Pujolle
Laboratoire Informatique de Paris 6 (LIP6)
Université de Paris 6
8, rue du Capitaine Scott, 75015 Paris, France
Email: daheb@rp.lip6.fr

Wissam Fawaz, Ken Chen
Laboratoire de Traitement et Transport de l'Information (L2TI)
Université de Paris 13
99, av. Jean-Baptiste Clément, 93430 Villetaneuse, France
Email: fawaz@rp.lip6.fr

***Abstract*— Optical networks are becoming an inevitable technology in the next generation Internet. Optical WDM networks require a control network for automatic and rapid configuration. A control plane is under standardization by the Internet Engineering Task Force (IETF) under the name Generalized Multi-Protocol Label Switching (GMPLS). However, this control plane is unaware of the specifications of the services to be enforced. This work presents a policy-based management architecture adapted to an emerging kind of optical networks called hybrid photonic networks. Contributing to network flexibility, we present extensions to the management plane of these transparent wavelength and switching capable networks that provide the means to leverage their inherent capabilities. The proposed management architecture would provide optical network operators with the possibility to provision optical services, defined through high level service contracts, in an efficient and dynamic way.**

Keywords: Hybrid Photonic Networks, Generalized Multiprotocol Label Switching (GMPLS), Policy-Based Management, Service Provisioning, Service Level Agreement.

I. INTRODUCTION

The explosive growth of data traffic, shaped primarily by the proliferation of Internet and Virtual Private Networks (VPNs), has created a demand for capacity that doubles every year. It is thus in this context that a transmission technology like optical Dense Wavelength Division Multiplexing (DWDM) becomes a key factor in accommodating the continuing expansion of demand. The revolutionary DWDM technology increases transmission capacity of fiber links by several orders of magnitude. Nonetheless, blindly augmenting transmission capacity has proven not being the long-term solution. In fact, the huge increase of capacity challenges the switching equipments managing the wavelengths conveyed along the fiber links. For instance, with regard to cost efficiency, although laying multiple fibers may help to reduce the transportation cost, yet it shifts the complexity and cost to the bottleneck switching and regeneration nodes. To cope with this limitation, the so-called hybrid photonic networks (HPN) are emerging. Yet, in order to enable rapid provisioning of these networks, a common control plane is needed. Such a control network is currently being standardized at the Internet Engineering Task Force (IETF) under the Common Control and Measurement Plane (CCAMP) working group as the Generalized Multiprotocol Label Switching (GMPLS). The GMPLS framework enables the fast and rapid provisioning of network services while ensuring certain characteristics. However, this control plane has no capability to derive service requirements that are deduced from service agreements. Owing to the crucial role of hybrid photonic networks in next-generation optical data transmission, which capability is augmented with the GMPLS control plane, this article proposes a policy-based approach for service provisioning adapted for such an environment. The main role of this architecture is to bridge the gap between the service level agreements and the control plane provisioning mechanisms.

Next section provides an overview and some details related to hybrid photonic networks and policy-based management. Section III, introduces the use of policy-based management in such networks by presenting a policy based architecture adapted to the specific needs of hybrid photonic networks. In section IV, a policy execution example is presented. Finally, section V concludes this paper.

II. POLICY-BASED MANAGEMENT OF GMPLS-ENABLED HYBRID PHOTONIC NETWORKS

In order to keep up with the incumbent challenges of traffic growth in a reasonable way, next-generation optical carrier networks are expected to support the increasing load by employing advanced transmission like DWDM, and intelligent switching technologies such as hybrid optical cross-connects. Indeed, hybrid cross-connects are constituted of a transparent waveband switching stage [1] and of a regenerative wavelength switching stage with a partial capacity with regard to the overall node throughput (Fig. 1). A waveband is formed by a set of wavelengths, and is either switched in the optical domain to another waveband, or dynamically directed to the wavelength switching stage where electronic processing is performed. The transition to the electronic stage is required to regenerate a wavelength, to aggregate traffic into it, or to switch it to another wavelength. Such a hybrid switching environment will heretofore be referred to as Hybrid Photonic Network (HPN) [2].

Configuring network elements in a HPN to provide a specific service requires flexible and simultaneous configuration of more than one network element. The concept of Policy-based Management addresses that problem and offers solutions. Policy-based management has been the subject of

Authorized licensed use limited to: Lebanese American University. Downloaded on January 25,2023 at 15:51:56 UTC from IEEE Xplore. Restrictions apply.

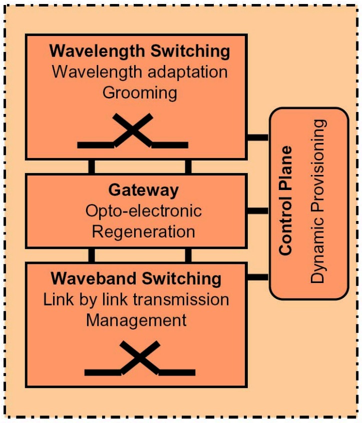

Fig. 1. Hybrid cross-connect architecture

extensive research over the last decade [3]. More recently, the Internet Engineering Tak Force (IETF) has been investigating Policy-based Networking as a means of managing multi-service networks with quality of service guarantees [4]–[6]. Policies are seen as a way to guide the behavior of a network or a distributed system through high level declarative directives. However, since the policy approach is a very general one and has to solve a number of issues simultaneously, it is useful to examine its application to GMPLS-enabled HPN [2], [7].

In order to backup our study with concrete business objectives related to the optical domain, the Optical Service Level Agreement (O-SLA), we defined in [8] was of great help. As a matter of fact, the O-SLA serving as a formal contract was intended to provide optical operators with guidelines on how to propose different optical services and service classes to their clients. To meet this purpose, different Service Level Specifications (SLS), embodied in the technical part of the O-SLA, were defined. Some policy rules needed for HPN management purpose were deduced from these O-SLSs and presented in previous works [9], [10].

During the provisioning of optical services in HPN, the management plane must operate in conjunction with the GMPLS control plane [7]. In fact, GMPLS provides a framework in which the well-known and proved MPLS paradigm is being extended to be a control plane for networks including both packet switching and circuit switching technologies. One of the merits of GMPLS stems from its ability to simplify the circuit provisioning process in optical networks. This simplification is realized through a suite of protocol extensions that are under standardization at the IETF.

GMPLS presents several functional blocks that are distributed along the different nodes in the network. The link state Interior Gateway Protocol, which can be either Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS) with optical specific extensions, is responsible for distributing information about topology, resource availability, and network status [11], [12]. This information is stored in a Traffic Engineering (TE) database. A constraint-based routing function acting as a path selector is used to compute routes for the desired Label Switched Paths (LSPs). This route calculation accounts for both the information collected in the TE database as well as the traffic requirements specified through the SLS parameters. Once the route has been computed, a signaling protocol such as Resource Reservation Protocol with TE extensions (RSVP-TE) is used for path activation [13].

In HPN, the internal topology abstraction consists of Traffic Engineering links (TE links) and the set of advertised Forwarding Adjacency, which are LSPs declared as links in the link state information base. The topology view will differ whether the computation of the explicit route is for a waveband-LSP or for a lambda-LSP. Following the LSP nesting principle, and in the considered HPN context, a waveband-LSP must be established before establishing the lambda-LSP to be nested in. The waveband coverage is a space-time dependent problem and is based on a predictive traffic analysis.

In order to provision and manage optical services in HPN, the management plane must operate in conjunction with the GMPLS control plane. This latter requires additional information to meet the operators expectations. In other words, the SLAs contracted between the operator and its clients provide the rules governing the interaction between the management and control planes. In this interaction scheme, the operator uses management functions to guide the control operations in an attempt to engineer the network according to business rules. As stated before, we base such an interaction scheme between the management and the control planes on network policies that we already presented in [9]. Hereafter, we will discuss the applicability and the framework used to control HPN via policies. A global policy architecture adapted to HPN is presented first, followed by a more detailed description of each of its components.

III. POLICY FRAMEWORK FOR HYBRID PHOTONIC NETWORKS

A. Global Policy Framework

The overall framework that has been adopted for the purpose of controlling HPN using policies is shown in Fig. 2. It is expected to assign a unique O-SLA Identifier (SLA-Id) for each client, once a service contract has been set with the operator.

The SLA-Id is needed to connect a client to the service being requested, as it is possible for one customer to contract several services and therefore O-SLAs with the same operator. In this regard, the SLA-Id would serve as a unique attribute based on which a specific service contract is identified by the management and control planes. Furthermore, achieving service provisioning in HPN necessitates an interaction between the management plane and the GMPLS control plane, this interaction is expected to be done via policy rules. As such the utility of the SLA-Id is augmented by the necessity to associate with each contracted O-SLA a number of policy rules whose role is to ensure the right enforcement of the service. Such an association is made possible through the use

Authorized licensed use limited to: Lebanese American University. Downloaded on January 25,2023 at 15:51:56 UTC from IEEE Xplore. Restrictions apply.

of the SLA-Id object, which can provide the link needed to index the different rules connected to a specific O-SLA.

Once the service contract has been settled and the SLA-Id attributed to the customer, the provisioning process of the identified service depends on two events. The first event consists in the receipt of the clients session request via the UNI interface [14], which entails a classical Call Admission Control (CAC) performed by the management plane. In the request, the message is conveying the SLA-Id object which is communicated to the management plane at the ingress in order to perform the CAC function on the identified user. The result of the admission control would be either to grant or deny the user access to network resources. The second event is the provisioning of the policy rules associated with the identified service request. Some of these policy rules are inferred from the service contract, and intend to make sure that the service is being deployed under good conditions. Others serve to implement and enforce operators objectives. Importantly, during this phase, the management plane downloads the rules to the ingress node.

With regard to chronological order of the above addressed events (represented by the arrows in Fig. 2), one event can precede the other according to the adopted provisioning strategy. In fact, two scenarios can be distinguished. In the first one, policies are downloaded a priori, which means policy rules pertaining to each O-SLA are provisioned by the management plane prior to customers request arrival (Arrow 5 precedes arrow 1 in Fig. 2). This could be useful for cases where operators tend to pre-provision lightpaths across their network, as the pre-establishment of lightpaths may be guided by the policies downloaded a priori. In the second one, policies are downloaded on a per session basis; it is only when a session request is received from the customers side that policy rules are provisioned into the ingress node, Fig. 2.

This is useful for cases where the operator establishes dynamic O-SLA with its clients. In other words, a dynamic O-SLA is a contract where a subset of SLS parameters can be changed easily over time. Service negotiation for a dynamic O-SLA is thus not performed manually but rather via a protocol transporting the varying SLS parameters. In the context of the present framework, service negotiation for a dynamic O-SLA could be done using the UNI interface which can be extended to transport objects related to SLS parameters. When doing so, the policies related to an O-SLA are not downloaded unless a service request is received from the customer specifying the desired values for the SLS parameters.

B. Technical Policy Framework

In the previous sub-section we proposed a general policy control framework that aims at controlling the service provisioning process in HPN networks. The main objective was to provide operators with guidelines on how to meet service requirements indicated through O-SLAs. However, within the proposed framework, the management plane was depicted as a black box and no additional details were provided with regard to the specific function of each entity comprising the management plane. Here, we present the technical specifications of the major building blocks involved during the management activity.

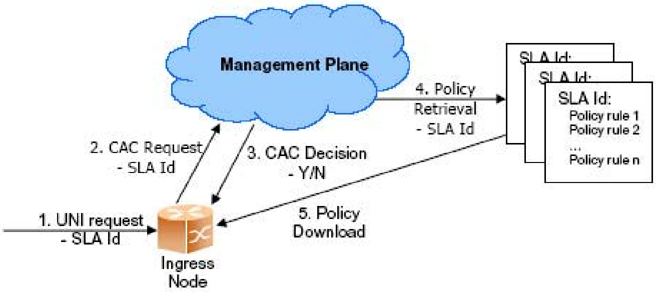

Fig. 2. Global policy enforcement framework

Fig. 3 depicts the various logical elements constituting the overall framework. Fulfilling the service requirements described by the O-SLAs is the main goal behind the definition of this framework. Policies are the facilitator used to achieve the fulfillment of service objectives and are derived from the O-SLAs. The translation, from high level objectives expressed through the O-SLA to policy rules like those presented in [9], has to be performed by a certain specific tool whose name is the policy manager, as defined in [3]. The policy manager will have as input the negotiated O-SLA then based on this input, it performs several validation tests and then translates the SLSs into policy rules. However, this translation process may be omitted. In such a case, the administrator will manually enter policy rules corresponding to each O-SLA together with the corresponding SLA-Id according to a specific file format (i.e., XML for example).

The Policy Decision Point (PDP), upon receipt of these policy rules, will store them in a policy repository (Policy DB on figure) which provides the central location that drives the service provisioning through the entire network. During subsequent operations, the PDP makes use of the SLA-Id to retrieve policy rules pertaining to a specific SLA. In harmony with specific events and triggers, the PDP communicates the policies to the Local Policy Decision Point (LPDP) which is responsible of the interaction with the Policy Enforcement Point (PEP) whose main task is to enforce policy rules. In this regard, the PEP refers (as shown in Fig. 4) to the optical node controller that runs the control plane protocols (such as OSPF-TE, RSVP-TE).

So each time there is a need to enforce a policy at the optical node level in order to meet some service objectives, the PEP makes a call in a way or in another to the LPDP which provides decisions associated with the designated service. The interaction between the PEP and LPDP will be detailed subsequently. It is this interaction that determines the interface which could exist between the management system and the GMPLS control plane [8].

Authorized licensed use limited to: Lebanese American University. Downloaded on January 25,2023 at 15:51:56 UTC from IEEE Xplore. Restrictions apply.

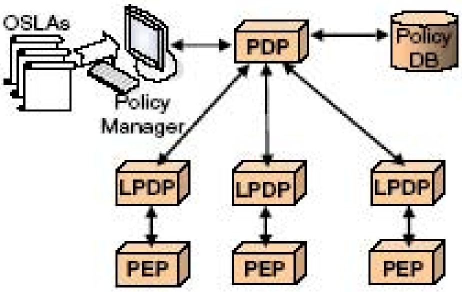

Fig. 3. Technical policy enforcement framework

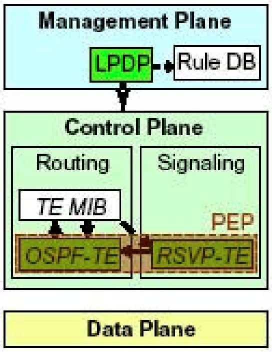

Fig. 4. Optical node functional blocks

C. *PDP and LPDP Specifications*

As stated previously, two service provisioning scenarios are to be distinguished. First, a pre-provisioning scenario in which policies are downloaded beforehand, that is prior to the receipt of customers request. Fig. 5, shows this Top-Down approach, in which the LPDP is responsible for triggering the service activation using the policies obtained from the PDP at the first beginning. Second, a post-provisioning scenario, in which only when service request is obtained from the clients side under acceptable conditions that the demanded service is activated. In this case, the LPDP could have the policies related to the requested service pre-installed in its local data base. However, if it is not the case the LPDP can inform the PDP about the need to obtain these policies. Fig. 6, shows the case where it is the client request that triggers the service provisioning. This case is more suitable for bandwidth on demand services.

In both scenarios, the LPDP has a restricted policy vision; in fact, it only knows the SLAs relevant to the specific optical node. As a result, it would not be possible for LPDP to take global decisions, i.e. network wide decisions. Therefore it is the PDP that performs the Call Admission Control procedure. The PDP and LPDP functions during the CAC procedure are shown in Fig. 6.

As a wrap up of all the above analysis, the PDP is supposed to fulfill the following functions:

- Store/retrieve policies in/from the policy repository;
- decide to whether accept or reject the client during the CAC procedure basing its judgment on the so-called CAC policies;
- and download the policies pertaining to a specific service towards the LPDP.

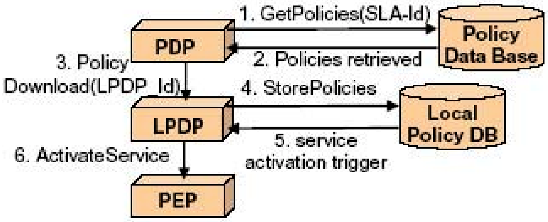

Fig. 5. Service activation scenario

On the other hand, the LPDP is responsible of the following functions:

- store/retrieve local policies in/from its local policy repository;
- and service activation through the interaction with the PEP.

D. *LPDP and PEP Interaction*

As shown if Fig. 4, the optical node control plane is composed of a routing block which computes paths and a signaling block which instantiate the paths computed. Both of these blocks need additional information provided by the management plane to meet service objectives specified by the O-SLAs or even operators objectives. So, it is clear that the optical node will interact with the management plane to complement its own operation. This interaction between the Network Element (NE) and the management plane is in fact more precisely a dialogue between the PEP running on this optical node and the LPDP. Actually, the PEP, which is a part of the management plane, is supposed to enforce policy rules that are downloaded from the LPDP. As a result it is clear that the PEP must be able to interact with the LPDP, via method call dialogue, in order to make use of the policy rules which are present at the LPDP level once they have been downloaded by the PDP.

IV. POLICY ENFORCEMENT EXAMPLES

In this section, we will be presenting two policy enforcement examples including the two possible management flow directions (Top-Down, and Down-Top). In more details, the first example will deal with a Top-Down approach where service activation is triggered at the management plane level, the flow of management messages is originated at the PDP and it is terminated at the PEP level. On the other hand, the second example presents a kind of situations where the management flow is triggered by the PEP and sent to the PDP for decision making purposes, meaning a Down-Top approach.

A. *Service activation example*

In order to clarify all of the above analysis, we present in this section an example related to the execution of a policy rule pertaining to service activation. The policy rule put into

Authorized licensed use limited to: Lebanese American University. Downloaded on January 25,2023 at 15:51:56 UTC from IEEE Xplore. Restrictions apply.

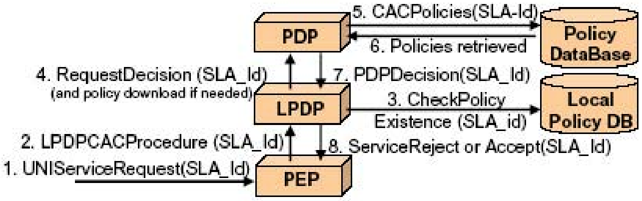

Fig. 6. Call admission control scenario

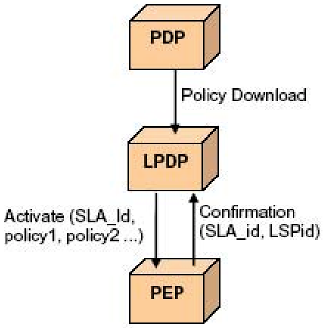

Fig. 7. A service activation policy rule scenario

action in this case is of the following form, expressed here in a high level language:

$$IF\ (Current\ date\ approaches\ service\ start\ date)\ THEN\ activate\ service\ in\ the\ network \tag{1}$$

If a pre-provisioning scenario is chosen for service activation, then the PDP will make use of similar directives in order to a priori download the policies into the LPDP. As a result, the LPDP would be able to activate the service prior to the service schedule stipulated in the O-SLA.

Whenever the current time approaches service schedule, the LPDP triggers service activation at the PEP level through the creation of a new LSP following the guidelines of policy rules. Which are communicated by the LPDP to the PEP through different function calls, since the LPDP and PEP are both implemented on the same NE.

A so-called Activate interface can be used by the LPDP, to interact with the NEs control plane making sure that service is provisioned according to the desired objectives (Fig. 7). For instance, at this stage, the Constrained Shortest Path First (CSPF) algorithm running at the NE level could take advantage of the routing policy rules of the service in order to calculate a route conforming to the service requirements. The policy rules (routing policies, for instance) are sent to the CSPF module as parameters of the activate method. When needed, the same interaction applies to the other component of the NEs control plane (RSVP-TE).

Once the route is pinned down, the PEP via a so-called Confirm method informs the LPDP about the LSP creation communicating an LSP-Id parameter which is used to facilitate possible future communication between the two entities (i.e., PEP and LPDP). So, the LPDP may use this LSP-Id to designate a specific LSP, especially if future modifications in terms of service requirements are needed to be taken into account through this same LSP.

With regard to the synchronization of the three boxes: PDP, LPDP and NE, since this policy rule is time aware, and in order to ensure a good execution of the policy rule, the three boxes must be tightly synchronized. Otherwise, synchronization mechanisms must be taken into account before the policy activation via the aforementioned method call.

B. Rerouting management example

In this subsection, we illustrate policies related to the concept of managing traffic trunks rerouting as presented in [10]. In fact, the O-SLA defined in the same reference included a parameter called Routing Stability that determines if optical traffic trunks can be rerouted or not. In the case in which it is agreed by the operator that the traffic trunks can be rerouted, this parameter specifies how often this will take place. Moreover, when an operator performs TE housekeeping within his network, he must make sure not to violate this parameter. As such, a policy rule must be put into action to prevent such a violation. In fact, what really happens is that when the load of a LSP drops below a certain threshold during a predetermined time period, it is more beneficial for the operator to reroute traffic in an attempt to optimize network resource usage. However, through our proposition in the O-SLA, we introduced the routing stability parameter to indicate how often this housekeeping may impact the clients traffic stability. The policy rule we proposed to accomplish the previously exposed objective has the following form:

$$\begin{aligned} IF(LSP\ load < threshold) AND(elapsed\ time = value)\\ AND(Number\ of\ rerouting < routing\ stability)\\ THEN reroute\ traffic\ and\ delete\ LSP \end{aligned} \tag{2}$$

Each LSP may aggregate several clients in the case where these clients request sub-lambdas connections. The more clients are aggregated in the LSP, the more load will be carried by the same LSP. Consequently, when the service of a client service is deactivated (service end time stipulated in the O-SLA is reached), the LSP traffic load will decrease. As such, the PEP detecting such an event sends to the LPDP the current LSP load, in order to determine whether the LSP rerouting is needed or not. If the LPDP has such an information (i.e., the corresponding policy rule) it can provide the PEP with the convenient decision to be enforced (whether to reroute the LSP or not). Otherwise, the LPDP requests this decision from the PDP side. The policy rule will provide the right decision to be enforced by the PEP, by comparing the LSP load to a certain threshold and making sure that rerouting will not violate the routing stability parameter. The decision is sent back to the PEP to trigger the rerouting process.

Authorized licensed use limited to: Lebanese American University. Downloaded on January 25,2023 at 15:51:56 UTC from IEEE Xplore. Restrictions apply.

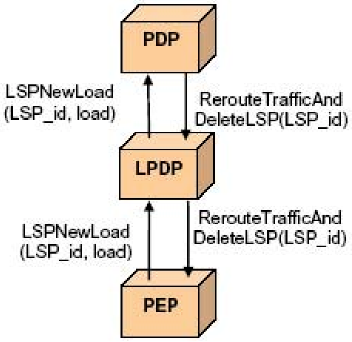

Fig. 8. A rerouting policy rule scenario

V. SUMMARY AND CONCLUSION

In a previous work we defined the different policy rules needed for the so called GMPLS-enabled HPN management, as a step towards stamping this kind of networks with the multi-service character. These policy rules took into account specifics of hybrid photonic networks with the help of a proprietary O-SLA. In this paper, a protocol independent policy control framework was presented as part of the overall solution. This framework instantiates the policy based management architecture and provides the technical specification of each component needed to accomplish the management purpose. Examples illustrating the different interaction schemes between these components and the GMPLS control plane were also discussed. In a future work, we will be concentrating on the implementation of this architecture, which will help us in the validation of the policy-based management architecture, providing the optical operators with the missing link needed to fulfill service objectives in optical networks in an automated manner.

REFERENCES

[1] R. Izmailov, S. Ganguly, T. Wang, Y. Suemura, Y. Maeno, and S. Araki, “Hybrid hierarchical optical networks,” *IEEE Communication Magazine*, vol. 40, no. 11, pp. 88–95, November 2002.

[2] M. Vigoureux, E. Dotaro, D. Papadimitriou, E. Oki, W. Imajuku, and N. Yamanaka, “Gmpls architectural considerations for (hybrid) photonic networks,” IETF,” Internet Draft, June 2002, draft-vigoureux-ccamp-gmpls-architecture-hpn-00.txt.

[3] D. C. Verma, *Policy-Based Networking: Architecture and Algorithms*. New Rides Publishing, 2001.

[4] P. Flegkas, P. Trimintzios, G. Pavlou, I. Andrikopoulos, and C. Cavalcanti, “On policy-based extensible hierarchical network management in qos-enabled ip networks,” in *International Workshop on Policies for Distributed Systems and Networks (Policy01)*. Springer-Verlag LNCS series, January 2001, p. 230.

[5] R. Yavatkar, D. Pendarakis, and R. Guerin, “A framework for policy-based admission control,” IETF,” Request for Comments: 2753, January 2000.

[6] D. Durham, J. Boyle, R. Cohen, S. Herzog, R. Rajan, and A. Sastry, “The cops (common open policy service) protocol,” IETF,” Request for Comments: 2748, January 2000.

[7] E. Mannie, “Generalized multi-protocol label switching (gmpls) architecture,” IETF,” Request for Comments: 3945, October 2004.

[8] W. Fawaz, B. Daheb, O. Audouin, M. Du-Pond, and G. Pujolle, “Service level agreement and provisioning in optical networks,” *IEEE Communication Magazine*, vol. 42, no. 1, pp. 36–43, January 2004.

[9] W. Fawaz, B. Daheb, O. Audouin, B. Berde, K. Chen, and G. Pujolle, “Policy-based provisioning in hybrid photonic networks,” in *International Conference on Parallel Processing Workshops (ICPPW’04)*. IEEE Computer Society Press, August 2004, pp. 407–414.

[10] B. Daheb, W. Fawaz, O. Audouin, B. Berde, K. Chen, and G. Pujolle, “Policy-based hybrid hierarchical optical networks,” in *International Conference On Broadband Networks (Broadnets04)*. IEEE Computer Society Press, October 2004, pp. 325–327.

[11] D. Katz, K. Kompella, and D. Yeung, “Traffic engineering (te) extensions to ospf version 2,” IETF,” Request for comments: 3630, September 2003.

[12] H. Smit and T. Li, “Intermediate system to intermediate system (is-is) extensions for traffic engineering (te), Request for Comments: 3784,” June 2004.

[13] L. Berger, “Generalized multi-protocol label switching (gmpls) signaling resource reservation protocol-traffic engineering (rsvp-te) extensions,” IETF,” Request for Comments: 3473, January 2003.

[14] O. Aboul-Magd, S. Ansorge, K. Arvind, K. Bala, S. Ballarte, A. Banerjee, R. Barry, D. Basak, G. Bernstein, R. Bradford, C. Brownmiller, Y. Cao, J. Drake, H.-M. Foisel, W. Goodson, G. Grammel, R. Graveman, E. Gray, R. Hartini, E. Mannie, R. Jain, L. Jia, J. Jones, S. Katukam, N. Komaee, J. Lang, M. Lazer, F. Liaw, Z.-W. Lin, L.-Z. Liu, and B. Mack-Crane, “User network interface (uni) 1.0 signaling specification, release 2,” Optical Internetworking Forum, Tech. Rep., February 2004.

Authorized licensed use limited to: Lebanese American University. Downloaded on January 25,2023 at 15:51:56 UTC from IEEE Xplore. Restrictions apply.