A Policy-based approach for User controlled Lightpath Provisioning

D.L. Truong[1,2], O. Cherkaoui[1] , H. Elbiaze[1], N. Rico[2], M. Aboulhamid[2]
[1]Dépt. d'Informatique, Université du Québec à Montréal
Case postale 8888, succursale Centre-ville
Montréal (Québec) Canada - H3C 3P8,
[2] Dépt. d'Informatique et Recherche Opérationnelle, Université de Montréal.
{truongtd,aboulhamid}@iro.umontreal.ca, {cherkaoui.omar, elbiaze.halima}@uqam.ca, rico@info.uqam.ca

Abstract

There is a growing need for e2e lightpaths for high volume data transferring applications such as GridFTP and SAN. They wish to dynamically deploy lightpaths over multiple management domains. Research, sponsored by Canarie Inc., is underway to enable "customer-empowered networks" and to experiment them with the Canadian research network CA*Net4. New signaling and control approaches using Web Services have been proposed. One difficulty is that each domain must retain the control of their optical network infrastructure and ensure proper allocation of optical resources. Hence, it is important that the signaling takes into account the management constraints imposed by the different domains.

This paper presents a policy-based approach for user-controlled lightpath provisioning. The work builds on the research around the new signaling approaches for realizing customer-empowered networks. We present an architecture based on Web Services allowing users or Grid applications to establish e2e lightpaths over multiple Autonomous Systems. To tackle the problem of admission control and to address the resource allocation issue, we developed policy restricted signaling which allows customers to reserve lightpaths over multiple domains while ensuring that management rules of each domain are enforced. The signaling has been implemented and the experiment on a small network composed of Cisco equipment proves the viability of our approach.

Keywords

User-controlled lightpath, Policy management, Inter-domain network management, Optical network, Resource reservation signaling, Web Services.

1. Introduction

The Ethernet over Sonet (EoS) technology has been integrated in various SONET switches, such as Cisco's ONS 15454 [11]. It allows carrying application data in Ethernet frames over Optical Transport Network (OTN) by encapsulating them within SONET frames. Two end hosts can be connected directly through highspeed Ethernet/EoS circuits. There is a growing need of e2e lightpaths for applications that require the transfer of high volumes of data such as GridFTP and Storage Area Networks (SAN). Users or applications need to flexibly control the provisioning of lightpaths across multiple independent domains because they are in a better position than the providers to choose and manage lightpaths adapted to their needs, for example in order to build optical VPNs. This has led to a new network paradigm called customer-controlled networks. In the paradigm, customers receive wavelengths or optical channels from a number of suppliers and control them independently, i.e. to establish and tear down lightpaths according to their need. Traditional approaches do not allow users to participate in the management.

Research, sponsored by Canarie Inc., is underway to enable those "customer-empowered networks" and to experiment with the CA*Net4. The goal is to give users or applications the ability to dynamically request network resources so to provide the users with the flexibility to develop network-based applications which require substantial network resources. The research resulted in a new signaling approach, which uses WebServices, called "User-Controlled Lightpath Provisioning" [1]. It allows customers or Grid applications to establish e2e lightpaths across multiple Autonomous Systems (ASs).

However, the proposed approach (see section 2) has two main drawbacks. First, no domain management is taken into account in resource allocation. It is important that the signaling approach takes into account the management constraints imposed by the different ASs (or domains). Otherwise, the user's lightpath configuration of optical resources can conflict with domain management policies and create undesirable effects. In this context, we believe that carriers and service providers should not leave the control of their lightpaths and network devices to customers without domain management and user discrimination on resource utilization. The use of policies will allow domains to flexibly define sophisticated rules that guarantee the domain management coherence.

The second drawback relates to the way e2e lightpaths are created and made available to customers. E2e lightpaths are searched, setup and then used right after. If the search fails, the customer will not obtain the lightpath and does not have time to react to the undesirable circumstance. On the other hand, if users could reserve lightpaths in advance, they are guaranteed to have them for a required time period.

Our policy-based solution (called UQAM Provisioning of Lightpath Application or UPLA), which is based on Web Services [3-4], addresses these weaknesses. The key ideas to overcome the above problems are i) to introduce an admission control mechanism governed by policies to regulate the customer intervention to the lightpath provisioning, and ii) to define a lightpath reservation mechanism that allows users to reserve e2e lightpaths in advance under the condition that admission rules of every network domains are respected. We use policies to moderate those admission rules.

The remainder of this paper is organized as follows. Section 2 summarizes related work including the Canarie's User-controlled Lightpath Provisioning Architecture (UCLP) as well as the GMPLS and the ASON lightpath provisioning model. Section 3 describes design approaches and the UPLA system architecture. Section

4 explains how the lightpath reservation and the lightpath deployment is realized using the proposed policy-based signaling approach. In Section 5, we present results of experiments performed on a small network which composed of Cisco devices. Finally, concluding remarks are given in Section 6.

2. Related works

The user-controlled aspect of our signaling is inspired from the UCLP. There are other well-known signaling approaches such as GMPLS and ASON, but we did not use them directly because of their limitation. This section reviews all three approaches.

2.1. User-controlled Lightpath Provisioning (UCLP)

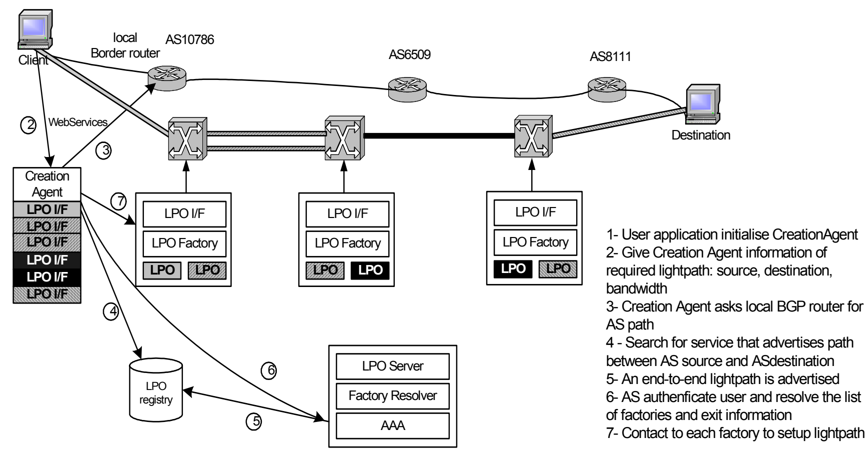

Figure 1: A simple case of lightpath set-up (see [1])

Canarie Inc. founded the Canadian network CA*net 3; the world's first national optical Internet research and education network. CA*net4, which builds on CA*net3, covers the Canadian territory and is composed of multiple ASs. Canarie proposed a user-controlled lightpath provisioning architecture that allows customers or Grid applications to establish e2e lightpaths over multiple ASs. Administrators of the ASs build up in advance short lightpaths (with pre-configured connection) which traverse one or more ASes and put them into a common lightpath repository called LPO Registry (LightPath Object Registry). The signaling for the establishment of an e2e lightpath consists of first looking into the LPO repository for short lightpaths, and then concatenating them to create the e2e lightpath. Figure 1 illustrates the different steps. To limit the search space, the Border Gateway Protocol (BGP) ASs path is referred to as the route of ASs across which the e2e lightpath should travel. In each AS, there are some dedicated modules for managing lightpaths and for controlling the set-up and concatenation. The unused e2e lightpaths may be advertised in the same manner, and other users can use them to build their own lightpaths.

2.2. Generalized Multi-protocol Label Switching (GMPLS)

GMPLS is proposed by IETF to switch the IP packets over a core optical network. A hierarchy of heterogeneous connection needs to be created with fiber on the bottom, then waveband, individual lambda, SONET/SDH tributaries, and packet switch capable connection at the top. GMPLS supports peer and overlay models [5,6]. The peer model is suitable for intra-domain problems because it is assumed that the edge nodes have the complete view of the core nodes' topology. The overlay model considers a non-packet based network as an AS and an AS in turn can be divided into sub-domains. An edge node of an AS does not need to be aware of the routing protocol used by the core nodes. The routing information exchanges between ASs are done via BGP-4.

2.3. Automated Switch Optical Network (ASON)

ASON [7] of ITU-T focuses on the inter-domain problem. ITU-T defines the UNI (User Network Interface) for the edge interface between networks and users or applications and NNI (Network Network Interface) for the interaction between network domains. Inside a network domain GMPLS is used.

However, both GMPLS and ASON are not customer-controlled provisioning. Except the connection set up, customers cannot exercise further on lightpaths like modification or partitioning as in UCLP. Moreover, GMPLS and ASON do not support advanced reservations as we will propose. Their reservations are for immediate use. In addition, both of them do not mention how domain management can be integrated in the e2e connection provisioning. Therefore, they have the same weaknesses as UCLP.

We employ the overlay model, which is similar to GMPLS or ASON and inspired by the signaling of UCLP. But different from these three architectures, we provide an architecture for intra-domain resource management and e2e lightpath provisioning which respects this management. The GMPLS or ASON model can be used inside domains in parallel with the domain resource management block as intra-domain solution. Similar to UCLP we are not interested in the intra-domain problem and assume that the domain reduces into a cross-connect.

3. UPLA layer model

In this section, we present our solution for user-controlled lightpath provisioning.

3.1. Design approaches

Our work was motivated by signaling approaches developed in Canarie's proposal. The following list describes the additional elements in our design:

- Maximizing domain autonomy: Every domain may accept or refuse a lightpath request according to its pre-defined management rules.
- Advanced reservation: Our lightpath reservation does not include the lightpath set up. There are two separated phases: i) searching for resources and marking them as reserved for the requested time period (called reservation phase) once

the reservation request is received, and ii) setting up resources in order to make them operational. This phase starts automatically by the system at the beginning of reserved time period (named setting up phase).

3.2. Architecture Overview

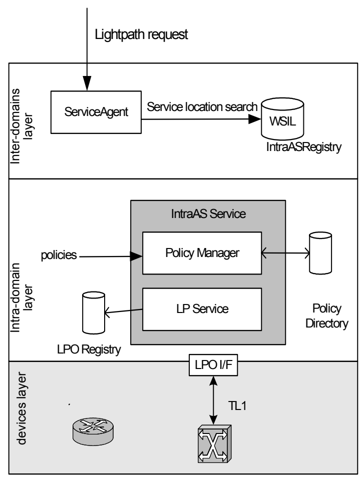

Figure 2: UPLA layer model

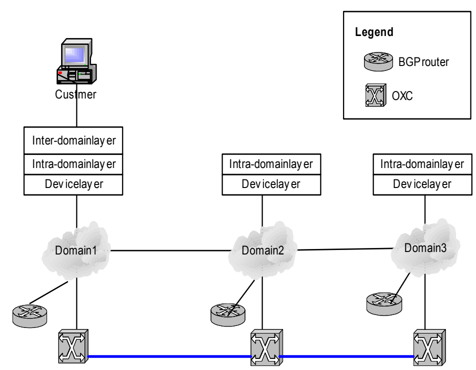

Figure 3: UPLA installation in multi-domain environment

UPLA is built for optical multi-domain networks composed of SONET/SDH switches. An e2e lightpath is a concatenation of a number of interconnected *single-domain lightpaths* of the same bandwidth, so called fragments. A single-domain lightpath starts from an interface of the ingress cross-connect of a domain and terminates on an interface of the egress cross-connect of the same domain. The later connects permanently with an ingress interface of the neighboring domain by fiber. The e2e reservation consists of a set of single-domain lightpath reservations and similarly, the e2e lightpath set-up consists of a set of single-domain lightpath set-ups.

The QoS guarantee for an e2e lightpath is naturally satisfied when all essential single-domain lightpaths are reserved because in optical TDM (Time Division Multiplexing) networks, bandwidth is granted by allocating a number of dedicated time slots in a wavelength. There is no bandwidth declining during the use as in packet switched network.

The architecture encompasses the following layers (see Figure 2 and 3):

- The Inter-domain layer is responsible for the reservation of e2e lightpaths by requesting single-domain lightpaths from domains and concatenating them.
- The Intra-domain layer is in charge of the lightpath reservation and setup inside a domain. It is also responsible for checking the conformance of lightpath requests with the domain admission rules.

- The device layer acts directly on the domain network devices. This layer allows the upper layers to be independent from the operating system of the network devices.

The interactions between layers or domains are realized by using Web Services.

4. Architecture description

This section presents the different layers of UPLA. We describe in detail the admission control mechanism and the lightpath reservation over multiple domains.

4.1. Inter-domain layer

Because an e2e lightpath is an aggregation of multiple single-domain lightpaths, lightpath brokers are used to collect lightpath fragments from different domains. For each domain, one lightpath broker, called Service Agent, is used to handle the e2e lightpath requests which originate from that domain. A global directory called IntraASRegistry provides ServiceAgents the service locations of every domain. The two important tasks of the ServiceAgent are to search and to reserve resources.

Resource searching mechanism

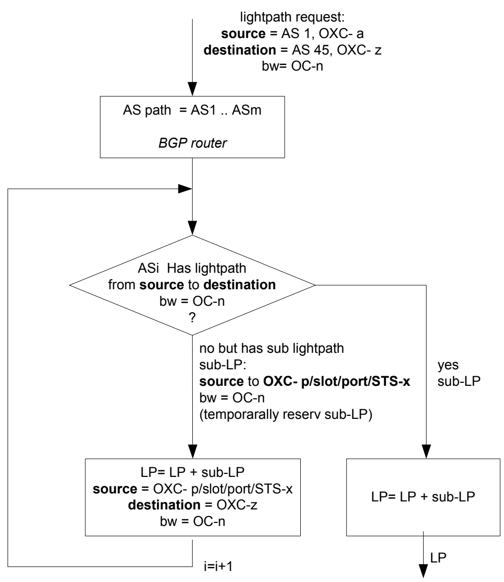

Figure 4: Lightpath searching algorithm

In the Canarie approach, a customer builds his e2e lightpath by selecting some pre-established lightpaths from a repository and then concatenating them. In the UPLA, no lightpath is pre-established to wait for being selected by the customer. AS by AS, potential single-domain lightpaths are searched, reserved and logically concatenated to obtain the e2e lightpath. The ServiceAgent is the principal actor of this process. Similar to ULCP, once an e2e lightpath request is received, the ServiceAgent starts by asking the local BGP router for the AS path from the source to the destination. The e2e lightpath is then formed by accretion as each of those ASes is probed in turn. Each AS will be asked for a potential single-domain lightpath that would extend to the destination. The lightpath may reach only midway to the destination if several ASes have to be involved. The next single-domain lightpath will start from the exit interface of the previous single-domain lightpath specified by cross-connect/slot/port/channel. This process continues through the different ASs until the destination is reached. These lightpath fragments are finally combined. Figure 4 illustrates this process.

If an intermediate AS does not have an available fragment lightpath, the search fails. However this search mechanism does not include back tracking because we assume that BGP router gives the correct inter-domain optical routing information. In fact, a conventional BGP router provides information about the connectivity of ASs but not about the availability of optical interfaces. Therefore, to strengthen the searching mechanism, OBGP (Optical BGP) [8] or other optical traffic engineering techniques should be introduced to routers in order to provide them the correct optical resource availability.

Reservation process

In order to reserve the e2e lightpath we have two reservation choices:

1) During the resource searching process, at each iteration, the found single-domain lightpaths are reserved immediately regardless of the final searching result. Note that the search may fail as mentioned in the previous section.
2) After identifying all essential single-domain lightpaths, they are reserved.

In the first case, we risk to reserve a single-domain lightpath for an e2e reservation which finally fails. Certainly, after being aware of the failure all reserved single-domain lightpaths will be released but they were unavailable in the interim. Therefore, another reservation request for the same single-domain lightpath at that time will be refused. This refusal is inappropriate because the single-domain lightpath is eventually free. In fact, the probability of this scenario may be considerable because the time needed to seek out resources for e2e lightpath increases. To avoid this situation, a single-domain lightpath should not be definitely reserved until all essential fragments for the e2e lightpath are identified.

In the second case, there is a time gap between the moment when a lightpath fragment is identified and when it is reserved. There is a risk that a reservation process reserves a lightpath fragment which is identified but has not yet been reserved by another reservation process. As a result, the reservation of the later process will fail because an identified lightpath fragment has been taken. The reason is that the two requests ask for the same critical lightpath fragment. In this case, the most important request should be accepted but not the faster one. The privilege level of requestors can be used to determine the important degree of their requests.

We propose then a two-steps reservation:

- Temporary reservation: During the resource searching process, all single-domain lightpaths proposed by ASs are immediately reserved temporarily. The resource under this reservation is available only for more privileged subsequent lightpath request and is unavailable for less privileged subsequent lightpath request.
- Reservation confirmation: Once the ServiceAgent has successfully assembled the needed lightpath fragments into an e2e lightpath, it confirms the ASes that the reserved fragments will be used. While many requests can temporarily reserve the same resource, only one confirmation can be accepted. The "winner" is the most privileged request among those requests which reserve the same lightpath fragment (concurrent reservation control). The resources are then committed and no longer available to any other lightpath requests.

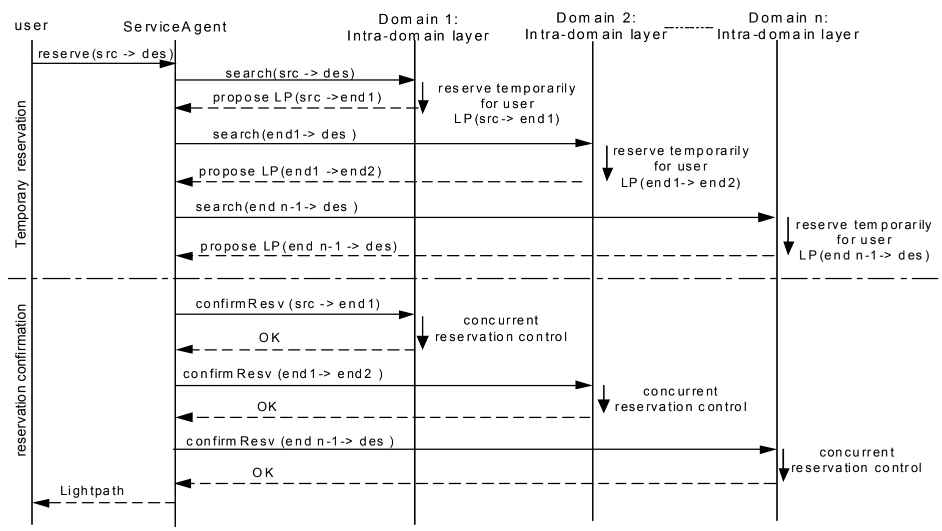

Figure 5: Inter-domains searching and reservation process.

The advantage of this reservation process is that a temporarily reserved fragment remains available for other more privileged users until it has been definitely reserved in the second step. Figure 5 highlights this process. In case the reservation fails in any domain, the ServiceAgent asks the precedent domains to cancel the confirm reservation. For the sake of simplification, we omitted to show the information forwarded between entities, such as the requested bandwidth, the time period, and the source/destination addresses as well as the exception cases.

4.2. Intra-domain layer

The Intra-domain layer is responsible for searching, reserving and setting-up single-domain lightpaths or lightpath fragments on behalf of customers and under the control of the domain management rules.

The intra-domain layer has 3 main components: LPServer, PolicyManager, and Intra-ASService which orchestrates the LPServer and PolicyManager. The LPServer manages the domain's optical resources consisting of lightpaths and interfaces (slot, port, STS channel). It controls device layer to set-up the reserved lightpath fragment.

The PolicyManager refers to the entire policy-based system with the Policy Decision, Policy Editor, Policy Directory, Policy Enforcement Tools (PET) modules. The PolicyEditor provides the domain administrator tools to define domain management rules and to store them in the Policy Directory. The Policy Decision module decides to accept or refuse a request for the domain's lightpath fragment according to those rules. Finally, PET triggers the setting up of a reserved lightpath when the reserved time period arrives.

Policy utilization in admission and resource reservation control

The concepts of policy and policy-based management are not new; we refer interested readers to [9-11] for details about these concepts. We explain here only how policies are used in UPLA.

An admission rule can be represented by a policy where the policy condition identifies user/application and the policy action assigns a privilege level to the user/application. For example, the following rule assigns to the *UQAM (Univ. of Quebec at Montreal)* user the privilege level 3 when he uses lightpath service:

if (user="UQAM") and (service=LPService) then (priority=3)

The greater the privilege level, the more privilege the user has. Note that this is a very simple type of management rule. Other more sophisticated rules can also be defined under the form of one or multiple policies.

Resource reservation can also be understood as a short-term management policy rule. The time conditions are the time intervals during which the lightpath must be in operation. For example, the request of the *UQAM* user for a lightpath number 108 from 8:00 to 17:00 is represented as:

(8:00 - 17:00) if (user="UQAM") and (LPO-ID=108) then "deploy lightpath"

where the LPO-ID is the identifier of the requested lightpath. For the sake of simplification, we replace lightpath setting-up call with many parameters by *"deploy lightpath"*.

Single-domain lightpath reservation

Similar to the Inter-domain layer, single-domain lightpath reservation consists of searching and reserving resources. In addition, the reservation includes the user admission verification. Figure 6 illustrates the process.

Once the Intra-AS Service has received a reservation request from the ServiceAgent (1), it asks the Policy Manager to check the user's right to use the lightpath service (2). If the Policy Manager responses positively (3), the Intra-ASService demands the LPServer to search for the requested single-domain lightpath (4). The LPServer then proposes one free single-domain lightpath (5,6). The Policy Manager is asked to make a temporary reservation on the fragment (7). Before reserving it, the Policy Manager verifies whether there is any pending reservation on the resources used by this fragment (8, 9). If not, the Policy Manager reserves the fragment temporarily (11, 12). Otherwise, the lightpath fragment will not be reserved unless the requesting user has higher priority than the owner of older temporary reservations (10).

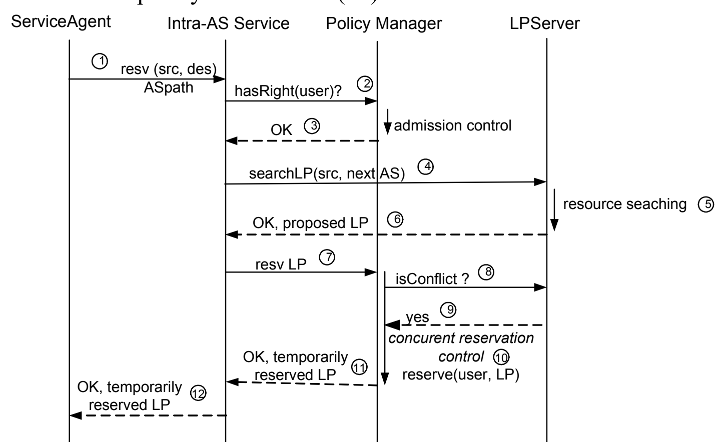

Figure 6: Single-domain lightpath searching and temporary reservation process

When a ServiceAgent confirms a reservation, the Policy Manager needs only to register the reservation and request the LPServer to mark the resources as busy.

Resource seeking for single-domain lightpaths

In step 5 Figure 6, the LPServer searches for resources and proposes a single-domain lightpath. The LPServer attempts to find a potential connection from a given ingress interface of the source domain (I1 cross-connect 1, figure 7) to any ingress interface of the destination domain (cross-connect 2). The connection is composed of: 1) free bandwidth that is requested for e2e lightpath over a link connecting the source and the destination cross-connect (I2-I3 connection), and 2) the cross-connection between the ingress and the egress interface of the source cross-connect (I1-I2 connection). The LPServer must look for such 3 interfaces and reserve the two of its own domain (I1, I2) in order to reserve the single-domain lightpath under consideration. The next single-domain lightpath for the e2e lightpath will continue from the I3 interface of the neighboring domain.

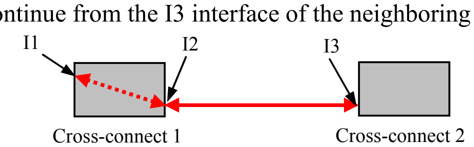

Figure 7: Single-domain lightpath (I1-I3)

Single-domain lightpath setup

In contrast to the lightpath reservation, the e2e lightpath set-up process is realized by independent set-ups of constituent single-domain lightpaths without cooperation between domains or user intervention. As lightpath fragments are already reserved separately in domains, the PET of the Policy Manager needs only to browse periodically the list of reservations to find the reservation whose starting time is approaching. Then the LPServer commands the device layer to make the essential cross-connection (I1-I2).

4.3. Device layer

The device layer consists principally of the LPO I/F, which talks directly to the cross-connect in order to create or break cross-connections. The LPServer uses these operations to setup or release single-domain lightpaths.

5. Experimentation and analysis

The UPLA has been implemented. Different experiments were conducted on a small network made of three domains named AS1, AS2 and AS3 (Figure 8). The cross-connects are CISCO ONS 15454 [12-14] and TL1 [15] is used to communicate with them. The ONSs are physically connected by OC-48 links between AS1-AS2, AS2-AS3. All the requests are for lightpaths from AS1 to AS3.

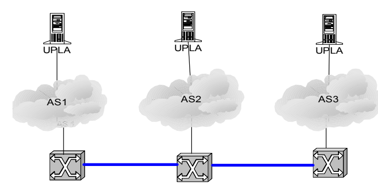

Figure 8: Experimental architecture

5.1. Lightpath reservation time analysis

The lightpath reservation process includes: asking the BGP router for the AS path, looking for service locations of each AS involved, searching and temporarily reserving a single-domain lightpath fragment in each domain and confirming these reservations. The time for asking a BGP router for a path or to look for service locations is small, so we did not include them in the analysis. We sent a sequence of 50 random lightpath requests over 3 domains. Table 1 presents the time taken by each signaling phase in millisecond.

Operations	**Average execution time (ms)**	**Minimum execution time (ms)**	**Maximum execution time (ms)**
E2e Reservation	17757	9132	25917
Search & temporary reservation in each domain (*)	3358	1623	7086
Reservation confirmation in each domain (**)	2475	777	5884
2 phases reservations time per domain (*) + (**)	5833	2400	12970

Table 1: Average execution time

Because the processing in each domain is almost identical, the e2e signaling time is proportional to the number of traversed domains. Thus, the *N*-domain lightpath reservation signaling time is:

$$Signaling\_time \approx N \times reservation\_time\_per\_domain \quad (1)$$

where *N* is the number of domains that the lightpath traverses; and the reservation time per domain is the time needed to reserve a lightpath fragment in a domain.
According to our experiment, the reservation time per domain is about 5.833s; therefore the e2e signaling time is:

$$Signaling\_time = N \times 5.833 \text{ (seconds)} \quad (2)$$

There are several factors causing this increasing signaling time. The first factor is the non-optimization of our current implementation, mostly in accessing the policy directory, the reservation base and the inventory. With the growth of the number of reservations, this factor has a bigger impact on signaling time. The use of WebServices with XML message processing is the second factor.
Since lightpaths are reserved for a future but not immediate utilization as in GMPLS or ASON, the signaling time is not a strict constraint. In the case of CA*net4 with around 12 management domains, the signaling time is acceptable. It shows that the approach is viable for national-scale networks.

5.2. Concurrent reservation control analysis

This section shows the experimentation how UPLA behaves in the case of multiple requests for the same resource. We create users corresponding to 3 privilege levels P1, P2, P3 with the following privilege order: P1> P2 > P3. The users request several times the same lightpath for the same time period. Note that shorter time delays between two consecutive requests increase the probability that there are conflicts in resource reservations. With a delay of 40s or higher, no conflict can happen because the delay is greater than the highest e2e signaling time, so every precedent reservation is completed before the coming if the new one.
Figure 9 presents the percentage of accepted requests for each user P1, P2, P3 when delay between two successive requests varies from the 5s to 80s. We note the conformance of these percentages with the privilege order of three users.
Figure 10 shows the percentage of requests that are refused in the searching phases. Some requests are reasonably refused because they intend to reserve the same resources as a more privileged request which has temporarily reserved the resource (but has not yet confirmed it). The refusal order is P3<P2<P1 which proves that increasing priviledges reduce the probability of being refused.
Figure 11 illustrates the percentage of requests that are refused in the reservation confirmation phase. In this phase, a request is refused only if it has temporarily reserved resources which are later temporarily reserved by a more privileged request. When the former request confirms the reservation it will be refused because resources are kept for the more privileged request. The refusal percentage of P2 requests is a bit higher than P1 requests because most of the conflicting requests of the P1 user have already been refused in the searching phase consequently less of them are refused in this phase. Some requests of P3 privilege level are refused because they are conflicting with other requests of the same privilege level.
The three experimental results show that if there are concurrent requests for the same resources, resources are kept for more privileged requests and they are accepted with higher probability.

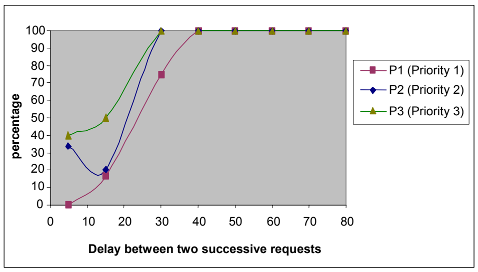

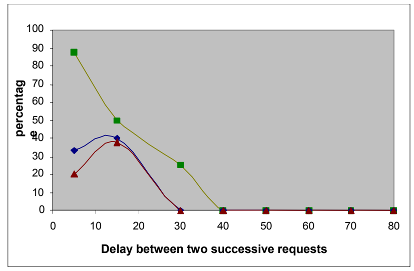

Figure 9: Percentage of accepted requests

Figure 10:Percentage refused requests in searching

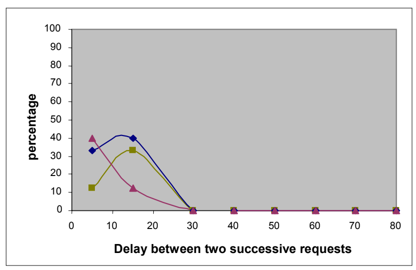

Figure 11: Percentage of refused requests in reservation confirmation

6. Conclusion and future work

In this paper, we presented a policy-based signaling architecture for dynamic provisioning of multi-domain lightpaths. The architecture allows each domain to manage its lightpaths independently through domain policies. The created lightpaths conform to these policies and are assigned to customers/applications based on their privilege level. Moreover, customers can reserve lightpaths in advance, so they will be guaranteed to obtain the needed lightpaths. The generalization of the experimentation shows that with the current implementation, the e2e signaling time is acceptable for national size multi-domain networks such as the CA*net4.

In UPLA access rights and privilege levels are defined as policies. Other more sophisticated domain management constraints can also be expressed in the same manner. Domain management is then simpler by relying on pre-defined management policies. This increases the flexibility of inter-domain cooperation and system scalability.

The UPLA architecture allows domains to join or leave an existing set of cooperating domains without any affect on global resource management since each domain manages its own resources and its management is guaranteed respected by

our signaling. The new domain needs only to register with the Intra-ASRegistry to make its presence known. The system scales up as new domains are added.
Further improvements on the implementation need to be made to decrease the signaling time per domain. Optimizations on resource allocation should be considered. Traffic engineering techniques can be taken into account for this objective. Those improvements are planned for the next phase of the project.

Acknowledgment

We express our gratitude to Mai Ly Vo Van and Daniel Puche from CISCO for their invaluable support throughout the project and for providing us access to ONS. The work could not have been completed without three students of our laboratory: Boubker Ghandour, Fabien Donato, and Khalid Ouyahya from École des Mines d'Alès EMA EERIE. We thank them for the implementation of UPLA and various precious ideas.

Reference

[1]. B. St-Arnaud, *"User controlled Lightpath Definition Document"*, http://www.canarie.ca/canet4/library/c4design/user_controlled_definition.ppt, Dec. 2002
[2]. CA*net4, *Canadian network*, www.cacarie.com
[3]. K. Ballinger, P. Brittenham, A. Malhotra, W. A. Nagy, S. Pharies, *"Specification: Web Services Inspection Language (WS-Inspection) 1.0"*, 2001
[4]. T. Modi, *"WSIL: Do we need another Web Services Specification? Explaining the difference between UDDI"*, http://www.webservicesarchitect.com/
[5]. E. Mannie, "*Generalized Multi-Protocol Label Switching Architecture"*, Internet Draft, May 2003
[6]. G. Swallow, J. Drake, H. Ishimatsu, Y. Rekhter, "*GMPLS UNI: RSVP Support for the Overlay Model"*, Internet draft, Oct. 2003
[7]. ITU-T, *"Recommendation G8080/Y.1304:Architecture for Automatic Switched Optical Network (ASON)"*
[8]. M. Blanchet, F. Parent, B. St-Arnaud, "*Optical BGP (OBGP): InterAS lightpath provisioning",* http://www.canarie.ca/canet4/obgp/obgp_ietf.doc, Mar. 2001.
[9]. H. Shai, *"Understanding Policy-Based Networking"*, John Wiley & Sons Inc., 2000.
[10]. J. Follows, D. Straeten, *"Application Driven Networking: Concepts and Architecture for Policy-Based Systems"*, IBM Redbook
[11]. B. Moore, E. Ellesson, J. Strassner, A. Westerinen, *"Policy Core Information Model -- Version 1 Specification"*, RFC 3060, Feb. 2001.
[12]. CISCO, *"Cisco ONS 15454 SONET Multiservice Provisioning Platform"*
[13]. CISCO, "*Cisco ONS 15454 User Documentation, Release 2.2.1"*
[14]. CISCO, *"Cisco Transport Manager Release 3.0 Operations Guide".*
[15]. CISCO, "*Cisco ONS 15454 TL1 Command Guide, Release 3.2*"