On the design of MPLS-ASON/GMPLS Interconnection Mechanisms

Luis Velasco, Ricardo Romeral, Fernando Agraz, Salvatore Spadaro, Jaume Comellas, Gabriel Junyent, and David Larrabeiti

Abstract— In this paper we propose a mechanism for connecting two or more MPLS islands belonging to the same MPLS domain through one ASON/GMPLS domain. It is based on the Overlay Model, where client and server networks do not exchange routing information. The interconnection is firstly done at the control plane level allowing the OSPF-TE flooding mechanism to advertise the existence of a link between two MPLS islands. Then, all MPLS routers in both parts of the MPLS domain know the complete network topology. Moreover, flooding messages advertising topological changes in one MPLS island of the domain are sent to other MPLS islands through these connections. No optical resources are used in the transport plane of the ASON/GMPLS network. New LSPs can be routed end-to-end triggering, if necessary, the establishment of LSPs in the ASON/GMPLS domain.

Index Terms— MPLS-GMPLS interworking, RSVP-TE, OSPF-TE, Overlay Model.

I. Introduction

MPLS and GMPLS interworking is a very active area for research and standardization. Within the IETF, the Common Control and Measurement Plane Working Group (CCAMP) [1] is leading a number of studies and extensions of existing protocols in order to achieve interworking between MPLS and ASON/GMPLS networks. Ref. [2] defines a general requirement list which any solution should carry out. The identified requirements are:

1. *End-to-end signaling*. MPLS signaling must be preserved through the GMPLS server network.
2. *Triggered establishment of GMPLS LSPs*. The set-up of end-to-end LSP in the client MPLS network may trigger the signaling of LSP in the GMPLS server network.
3. *Diverse path for end-to-end MPLS-TE LSPs*. The ability to establish protected end-to-end LSPs with diverse paths.
4. *Advertisement of MPLS-TE information via the GMPLS network*. The needed advertisement exchange between MPLS and GMPLS networks.
5. *Selective advertisement of MPLS-TE information via a Border Node*. The distribution of selective TE reachability information between MPLS and GMPLS networks.
6. *Interworking of MPLS-TE and GMPLS protection*. The support for MPLS Fast Reroute in the GMPLS domain [3].
7. *Independent failure recovery and reoptimization*. The independence of failure recovery and reoptimization mechanisms in the MPLS and GMPLS networks.
8. *Complexity and risks*. The introduced complexity should be as lower as possible.
9. *Scalability considerations*. The solution must be scalable with the number of GMPLS and MPLS nodes, the number of MPLS networks, and MPLS and GMPLS LSPs.
10. *Performance considerations*. The solution should be evaluated against the failure and restoration times and its impact over the control and transport planes due to the added overheads.
11. *Management considerations*. The needed management coordination between client and server networks.

In this paper, we propose a mechanism for connecting two or more MPLS islands belonging to the same MPLS domain through one ASON/GMPLS domain. We assume that MPLS and ASON/GMPLS network are managed by different operators; therefore, we are in a multi-domain scenario. The ASON/GMPLS network provides on-demand GMPLS LSPs to their client MPLS networks. The proposed solution will be evaluated according to the previous requirements.

However, in [2] the recommended solution architecture is based on the Border Peer Model. In this model, the interface between the client (MPLS) and the server (ASON/GMPLS) networks is the Peer Border Node. This node has full topology visibility of both networks. Routing information is not distributed from one to another network.

On the contrary, since we are in a multi-domain scenario with two different administrative networks, our solution is based on the Overlay Model [4]. This model implies a request/response protocol between client and server networks, the User-Network Interface (UNI) [5]. The key elements defined in this model are the Edge Node and the Core Node. Edge nodes belong to the MPLS domain and core nodes belong to the ASON/GMPLS domain and are administered by different network operators.

Manuscript received November 19, 2007. This work has been partially founded by Spanish Science Ministry through TEC-2005-08051-C03-02 “RINGING”, TEC 2006-27651-E “Redes Multinivel: IP sobre redes de transporte” and TEC2004-05622-C04-03 “CAPITAL” projects and by the EU through the project NoE e-Photon/ONe+.

LV, FA, SS, JC, and GJ authors are with the Optical Communications Group, Universitat Politècnica de Catalunya (UPC), Barcelona, Spain (e-mail: {luis.velasco, agraz, spadaro, comellas, junyent}@tsc.upc.edu).

RR and DL authors are with Dpt. of Telematics, Universidad Carlos III de Madrid (UC3M), Madrid, Spain (e-mail: {rromeral, dlarra}@it.uc3m.es).

X. Masip, S.Sánchez, J.Solé, A.Jukan (Eds.): ONDM 2008

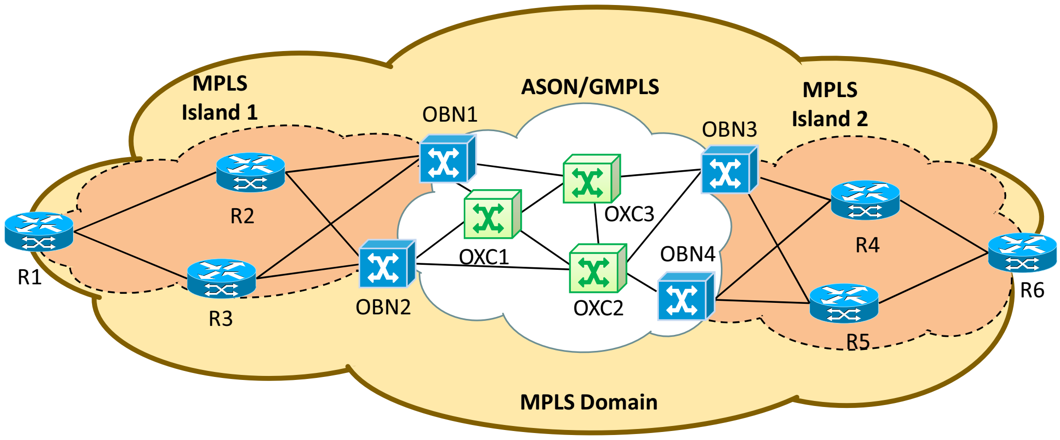

Fig. 1. Example of two MPLS islands connected through one ASON/GMPLS domain.

Edge nodes can be connected to several core nodes, in order to provide redundancy. However, directly connected edge and core nodes are closely related. We call this one-to-one relationship an Overlay Border Node, or simply a Border Node, in contrast to the Peer Border Node introduced in [2]. In the Overlay Model edge nodes and core nodes do not exchange routing information.

Three ways for signaling a LSP across domains have been defined: nesting or hierarchical, contiguous, and stitching [6]. However, contiguous and stitching options are only available when both domains are packet-based. Therefore, end-to-end MPLS LSPs are nested within another LSP crossing the ASON/GMPLS domain [7]. GMPLS LSPs created to support end-to-end MPLS LSPs are advertised as TE-links in the MPLS domain using a forwarding adjacency (FA), as defined in [7]. Then, MPLS LSP can be routed through these TE-links.

The remainder of the paper is organized as follows: Section II defines the network architecture. In Section III we detail our solution for interconnect MPLS islands through an ASON/GMPLS domain. Section IV describes the signaling process. In Section V we compare our solution against the requirements previously defined, identifying advantages and drawbacks. Finally, in Section VI we draw the main conclusions of this work.

II. Network Architecture

The main objective of this work is to design novel mechanisms that are needed to connect several MPLS islands through one ASON/GMPLS domain. An example of such a situation is shown in Fig. 1, where MPLS LSRs R1, R2 and R3 are in the MPLS island 1, and MPLS LSRs R4, R5 and R6 are in the MPLS island 2. Both MPLS islands are connected to the ASON/GMPLS domain through the Overlay Border Nodes (OBN) OBN1, OBN2, OBN3 and OBN4.

Border Nodes are a special kind of nodes belonging to both the MPLS network and the ASON/GMPLS network. Border Nodes are capable to manage both MPLS and GMPLS protocol sets and provide the required interworking. We define the Overlay Border Node functionality as two main collaborating functional blocks: the MPLS and the ASON/GMPLS blocks (Fig. 2). Henceforth we use the name Border Node to identify the Overlay Border Node here introduced.

The MPLS functional block uses the UNI interface to request GMPLS LSPs. In the GMPLS overlay model, MPLS functional block acts as an edge node and the Optical Connection Controller (OCC) sub-block acts as a core node for signaling purposes. The MPLS functional block is responsible for the following tasks:

- To establish, tear down, LSPs using resources in the existing TE-Links. RSVP-TE messages are sent to its Border Node neighbor through the existing TE-Link.
- To request new LSPs over the ASON/GMPLS domain using the UNI interface. GMPLS LSPs will become a TE-Link in the MPLS domain. GMPLS LSP requests are triggered by the signaling of MPLS LSPs, as we will detail in the next section.
- To propagate topology advertisements in the MPLS island which is connected to the rest of MPLS islands, sending OSPF-TE messages to its Border Node neighbors.

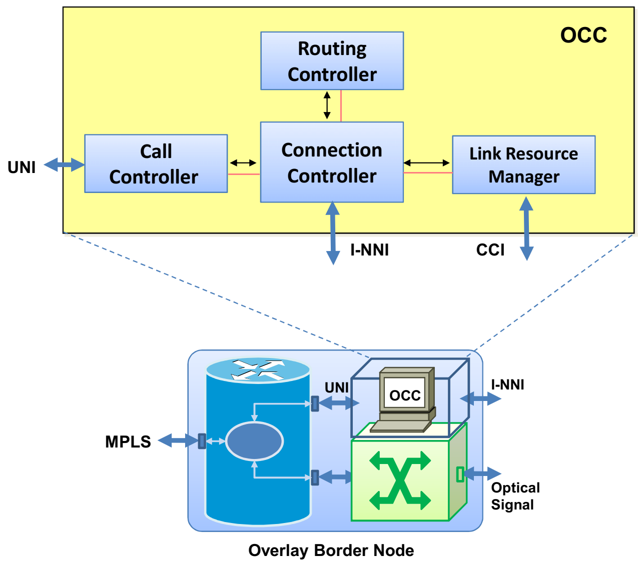

Fig. 2. Functional blocks of an Overlay Border Node.

The ASON/GMPLS functional block is divided into two sub-blocks: the OCC sub-block in the control plane and the Optical Cross-connect (OXC) sub-block in the transport plane.

The OCC sub-block contains the architectural components described in [9]. Specifically, it contains: 1) The Call Controller component, which is responsible for accept new call requests. LSP requests from the MPLS functional block arrive to this component through the UNI; 2) the Connection Controller component, which is responsible for connections set-up, modification and tear down. Communication between Call Controller and Connection Controller components is needed to requests new LSPs; 3) The Routing Controller component, which is responsible for routes computation; 4) The Link Resource Manager, which is responsible for manage link subnetwork points. This component uses the Connection Control Interface (CCI) to manage the OXC sub-block.

The OXC sub-block corresponds to an optical equipment with client interfaces (i.e. Gigabit Ethernet) and optical (DWDM, Dense Wavelength Division Multiplexing) interfaces with a number of wavelengths. The OXC block transforms client signals into wavelengths and multiplexes them into a DWDM signal.

The MPLS and the OXC functional blocks are connected through some of the client interfaces, with a specified bandwidth. MPLS data and related signaling use these interfaces.

III. MPLS-GMPLS INTERCONNECTION

In this section we analyze how to connect two or more MPLS islands of the same MPLS domain through one ASON/GMPLS domain. Different approaches can be considered:

- By effectively creating a lightpath at the deployment time of the Border Node. This means creating up to *n(n-1)/2* lightpaths to connect every Border Node with the remaining *n-1* Border Nodes in a full mesh network. This number grows sharply when the number of Border Nodes increases falling into unnecessary costs.
- By mapping MPLS protocols into GMPLS protocols. In this option MPLS messages between Border Nodes are encapsulated into their respective protocols in the GMLS domain, and are sent hop-by-hop through the GMPLS control plane until they finally arrive to the end Border Node. This option requires increasing the processing resources and the bandwidth in the GMPLS control plane.
- Using an IPsec tunnel through the GMPLS control plane to connect two Border Nodes, as shown in Fig. 3. This solution provides further advantages: firstly, it does not need neither mapping nor hop-by-hop processing for MPLS messages and completely separates MPLS and GMPLS control planes; secondly Border Nodes are connected by a single hop; and finally, MPLS control plane messages are transported encrypted through the GMPLS domain.

IPsec tunnels can be created at the deployment time of Border Nodes or can be added later. The objective of IPsec tunnels is to connect two islands of one MPLS domain and advertise this special link to the complete MPLS domain. As this special links exist only as entries of Traffic-Engineering databases in the control plane of MPLS routers, we call them as control plane links (cp-links). Fig. 4 shows an example of a MPLS domain connected through an ASON/GMPLS domain, where dotted lines represent cp-links. Cp-links have the following main characteristics:

- Cp-links are used to advertise the complete topology of one MPLS network. They connect two Border Nodes of the same MPLS domain. Cp-links are supported by IPsec tunnels through the GMPLS control plane. One Border Node uses IPsec tunnels to exchange OSPF-TE messages with its neighbor Borders Routers.
- The cp-links metric value should be higher than normal in order to avoid creating lightpaths in the ASON/GMPLS domain unnecessarily. Managing this value allow the MPLS network operator to reach high levels of grooming, minimizing the number of supporting lightpaths needed and therefore its costs.
- The available bandwidth of cp-links should be set to the bandwidth value of the interfaces available in the Border Node (those connecting the MPLS and the OXC blocks).
- One cp-link represents a number of potential future TE-links between the cp-link end points. On the contrary, if two Border Nodes are not connected through a cp-link, no TE-links will be created between them.

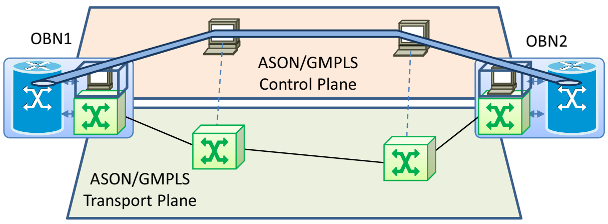

Fig. 3. IPsec tunnel between two Border Nodes created through the ASON/GMPLS control plane.

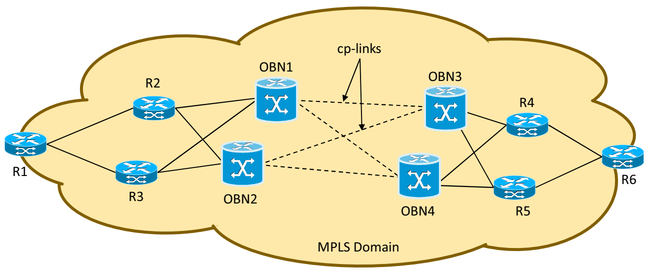

Fig. 4. Complete topology view at the MPLS Domain.

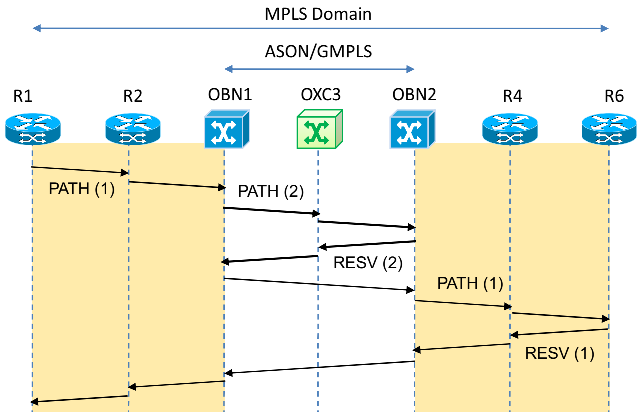

Fig. 5. RSVP-TE Path and Resv messages to create an end-to-end MPLS LSP. This signaling triggers the provisioning of a new GMPLS LSP.

- MPLS LSPs can be routed through cp-links. When a Border Node receives a RSVP-TE Path message for a LSP routed through a cp-link, it will create a new lightpath in the ASON/GMPLS domain with destination the remote Border Node for this cp-link. Once the lightpath has been created, the MPLS signaling can continue. This process is shown in Fig. 5.

After a new lightpath has been created in the ASON/GMPLS domain, it is advertised as a new TE-link in the MPLS domain using the OSPF-TE flooding mechanism. The TE-link metric value must be lower than that of the cp-link connecting the same Border Nodes –as MPLS domain does not know the number of hops in the ASON/GMPLS domain, this value has to be configured in base to other considerations. This way, new LSPs will be routed through existing TE-links instead of cp-links. Note that routing a LSP through a cp-link means creating a new lightpath in the ASON/GMPLS domain.

After the first TE-link has been created between two islands of a MPLS domain, new LSPs can be routed using this TE-link. Any new LSP created through a TE-link consumes part of the bandwidth available on this TE-link. Due to the lack of bandwidth, in some moment a LSP will be routed through a cp-link and, as a consequence, a new TE-link will be created. If the number of interfaces connecting the MPLS and the OXC functional blocks in the Border Node is not enough to create new TE-links, the available bandwidth of every cp-link ending in this Border Node must be set to 0. Advertising the new state of cp-links will prevent using them for routing new LSPs.

It is also possible that border nodes announce the links between them (cp-link and created TE-links) as a single one. When a new LSP use BW in one of the TE-links an OSPF-TE message is sent to all LSRs with the modified free BW. The advantage of this view is when a TE-link is sign-up or sign-down for all LSRs the unique change is the free BW of the (virtual) link, therefore is not necessary to execute the routing algorithm in all MPLS nodes. This was necessary when new links (and deleted links) was announced via OSPF.

Another improvement over the proposed solution is to divide the MPLS domain into different OSPF areas: One area per island and the area 0 or backbone containing all Border Nodes. In this way, the exchanged routing information between LSRs can be decreased when TE-link are created or deleted.

These two advances can help to improve the scalability in terms of control messages exchanged and CPU usage in the LSRs.

IV. MPLS-GMPLS SIGNALING

In the last section a new mechanism to interconnect MPLS islands through an ASON/GMPLS domain was defined. In this section we detail the MPLS-GMPLS signaling mechanism with the help of the network topology depicted in Fig. 1.

Let us imagine the case where the router R1 in the MPLS island 1 wants to establish a new LSP to R6 in MPLS island 2 to support a traffic demand. We assume source routing. Previously, OBN1-OBN3 cp-link has been created and advertised to all nodes in the MPLS domain. So, R1 creates a new RSVP session and compute the route as:

$$R1\text{->}R2\text{->}OBN1\text{->}OBN3\text{->}R4\text{->}R6 \quad (1)$$

In the case when a unique "virtual" link is used between Border Nodes the nodes specified in the ERO object must be set to "loose hops" [11]. If a multiarea OSPF domain is used then is possible to know which Nodes are Border Nodes (those where interconnection area nodes) marking as "loose hops" the GMPLS hop.

Then, R1 creates the RSVP-TE Path message for this session identifying the end points of the LSP in the SESSION and SENDER_TEMPLATE objects (R6 and R1), the explicit route in the ERO object, and information about the traffic to be sent through this LSP in the SENDER_TSPEC object. R1 sends the Path message downstream towards R2 which will process and forward it downstream to the next hop defined in the ERO object.

When the Path message arrives to OBN1, it gets, from the ERO object, OBN3 as the next hop. At this moment OBN1 identifies that OBN3 is connected though a cp-link. OBN1 has available local ports so decides create a new LSP through the ASON/GMPLS network to OBN3. Then, it creates a new RSVP-TE Path message identifying the end points (OBN1 and OBN3). As the LSP will be advertised as a forwarding adjacency, it includes a LSP_TUNNEL_INTERFACE_ID object in the path message with a valid local identifier for the forwarding adjacency [12]. To identify the local interface chosen, OBN1 includes an IF_ID RSVP_HOP object. No ERO object is needed in this case. OBN1 sends the Path message to the OCC of the OXC containing the end port associated with the local port chosen, through the UNI interface. Note that OCC, OXC and neighbor OBN IP addresses are configured at deployment time.

Upon receiving a Path message through the UNI interface, the Call Controller in the OCC will apply the local policy in order to accept or reject the new call. Once the call is accepted, the Connection Controller asks the Routing Controller to compute a route to the destination. If a route exists, it creates a new ERO object with the route, attaches it to the Path message and sends the message to the next hop in the route.

When the Path message arrives to the OCC in OBN3, it removes the ERO object from the incoming Path message and sends the message to the MPLS block destination through the UNI interface. After processing the Path message, the MPLS block in OBN3 will create a MPLS Resv message and it will include a LSP_TUNNEL_INTERFACE_ID object with a valid local identifier for the reverse end point of the forwarding adjacency. The Resv message is send upstream towards OBN1.

When the Resv message arrives to OBN1, and after the corresponding reservations have been performed, the LSP between OBN1 and OBN3 is available to be used. At this point, end-to-end R1->R6 LSP signaling can continue using the just created LSP to go through the ASON/GMPLS domain.

OBN1 (and also OBN3 if the LSP is bidirectional) can advertise the LSP between them in the MPLS domain as a new forwarding adjacency by means of the OSPF-TE flooding mechanism. New LSPs can be established between both MPLS islands by using this new TE-link for routing.

Forwarding adjacencies in the MPLS domain are removed, and their corresponding LSP in the ASON/GMPLS domain are torn-down, when the last MPLS LSP using them is torn-down. This may happen when the last end-to-end LSP using the forwarding adjacency is torn-down or when the signaling of the LSP which has triggered the establishment of the forwarding adjacency fails in the second MPLS island, and the corresponding PathErr message is received by OBN1.

V. SOLUTION EVALUATION

After the in deep description of our proposed solution to interconnect MPLS islands through an ASON/GMPLS domain presented in previous sections, in this final section we evaluate our solution against the requirements identified in [2]. Moreover, advantages and drawbacks of the solution are also identified.

1. *End-to-end signaling*. MPLS signaling is not accessed by the ASON/GMPLS since end-to-end MPLS LSP signaling is always sent over the transport plane of the ASON/GMPLS domain. When a PATH message arrives to a Border Node two situations can be considered: First, if a TE-link exists and has enough BW, the PATH message is sent to the next Border Node through the ASON/GMPL domain, and is received by the MPLS functional block in the destination Border Node; Second, if the Border Node has to create a new LSP through the ASON/GMPLS domain, a new and independent GMPLS signaling procedure, based on the UNI interface, start to sing up a new LSP. Then, when it is created the procedure is as previous case, the end-to-end MPLS signaling continues normally. Of course, this procedure introduces a small delay over the end-to-end LSP signaling.
2. *Triggered establishment of GMPLS LSPs*. When a PATH message arrive to a Border Node with a next hop in the ERO Object specifying a Border Node, the receiving Border Node can use an existing TE-link or create a new one and use it. This decision is transparent to the source LSR.
3. *Diverse path for end-to-end MPLS-TE LSPs*. Diverse path computation has not been studied here but if cp-links (and then TE-links too) use SRLGs [12] to indicate if shared resources are used and diversity inside de ASON/GMPLS domain exists, then source LSR can compute disjoint paths using OSPF information and sing up disjoint LSPs.
4. *Advertisement of MPLS-TE information via the GMPLS network*. The cp-links links are advertised to LSR of the MPLS domain via OSPF-TE, cost and available BW attributes may be used. Cp-links are advertised in an opaque way, i.e. without information about used resources in the ASON/GMPLS domain. cp-links are used to send routing control information between MPLS islands. These links ensure that control connectivity inter-islands is always available. Data links are dynamically created (see point 2).
5. *Selective advertisement of MPLS-TE information via a Border Node*. Obeying with interdomain exchange network information restriction, that is, the topology of a domain must be opaque to others, the view of MPLS islands of the GMPLS domain is just the links defined between different Border Nodes as a single MPLS link. Overlay Model used in this solution guarantee that not confidential information is sent to MPLS network from Border Nodes.
6. *Interworking of MPLS-TE and GMPLS protection*. As links in the ASON/GMPLS network is viewed as a single hop by LSR of the MPLS islands it is possible to find alternative paths to protect any ASON/GMPLS link or Border Node using Fast Reroute [3]. Fast Reroute can be used in all resources of the MPLS network. Inside the ASON/GMPLS network Fast Reroute can be used to protect OXC nodes.
7. *Independent failure recovery and reoptimization*. If a failure occurs in the ASON/GMPLS network, a local recovery do not affect to the end-to-end MPLS LSP. The label relationship Border Nodes remains with no changes. The output/input interfaces between MPLS and OXC blocks inside each border nodes also remain unaltered. Then, neither changes, fast reroute actions nor reoptimization inside the ASON/GMPLS network affect MPLS network.

   Changes over a LSP in the MPLS domain (fast reroute o reoptimization) do not affect to ASON/GMPLS network as this network is serving connections between Border Nodes. If due to some fast reroute o reoptimization action in the MPLS network, one LSP is rerouted from one to another TE-link between two Border Nodes, this action is performed without having any impact over the ASON/GMPLS domain.

8. *Complexity and risks*. The ASON/GMPLS network is seen as a set of links between Border Nodes by the MPLS domain. Thus, the added complexity to the MPLS domain is minimum. On the contrary, two requirements appear on the ASON/GMPLS control plane side: 1) The UNI interface by which MPLS network requests new LSPs, and 2) the need of connecting MPLS functional blocks in two Border Nodes through IPsec tunnels. Thus the added complexity is very low.
9. *Scalability considerations*. Only two scalability issues are involved: 1) IPsec tunnels. To create a full mesh network connecting n Border Nodes, $n(n-1)/2$ IPsec tunnels have to be created trough the ASON/GMPLS control plane. However, the IP traffic to be borne by the ASON/GMPLS control plane is limited to OSPF-TE messages; 2) OSPF-TE advertisements. Every change in cp-links and the related TE-links connecting MPLS islands have to be announced to all LSRs the MPLS network. This can be solved using OSPF areas, as explained on previous sections.
10. *Performance considerations*. Using fast reroute in both, MPLS and ASON/GMPLS networks the recovery times can be as short as the recovery time in a MPLS or ASON/GMPLS networks. The only exception is when a Border Node fails; in this case, a reroute over another TE-link in other Border Node must be performed and can be necessary to create new TE-links, if a backup does not exist, adding an extra delay to the recovery time.
11. *Management considerations*. With manually configured cp-links, the MPLS network works automatically. OSPF-TE offers the needed tool to connect two MPLS islands. With the islands connected through cp-links and the Border Nodes doing the ASON/GMPLS data connections, the network works fine. Moreover, MPLS signal protocols can be used to perform failure notifications.

VI. Conclusions

In this paper we have presented a proposal to connect several MPLS islands of the same domain through an ASON/GMPLS domain. This solution allows to sign-up end-to-end MPLS LSPs using a third party network (ASON/GMPLS network).

We define the cp-links as links connecting Border Nodes in different MPLS islands. Cp-links are used to define reachibility among Border Nodes. Cp-links are based on the establishment of IPsec tunnels trough the ASON/GMPLS control plane. Border Nodes use IPsec tunnels to announce MPLS network routing and reachability information among MPLS islands. Cp-links can be established at any time after the deployment of Border Nodes.

TE-links are ASON/GMPLS connection between Border Nodes to transport data packets from one island to another one. When a PATH message signaling a MPLS LSP arrive to a Border Node this use an existing TE-link or if is full sign-up a new one and use it. All this process is transparent for MPLS network. Note that MPLS LSP signaling is transported between a pair of MPLS islands, trough the existing TE-link.

The proposed solution is in accordance with the requirements defined in [2]. In particular, the MPLS signaling is performed end-to-end in a transparent way. MPLS signaling triggered the signaling, if is necessary, of a new TE-link on ASON/GMPLS network.

The topological network information is hidden between MPLS and ASON/GMPLS networks. This point is very important because the internal topology must not be shared between two different operators.

With a fast recovery (local recovery) strategy the recovery of an end-to-end MPLS LSP can be performed individually in the different networks and with similar detection and recovery times of an only-one-operator network.

Finally, our solution is completely scalable; cp-links interconnect Border Nodes and only are used for OSPF-TE messages. The maximum number of possible cp-links is $n(n-1)/2$, where n is the number of Border Nodes. This makes a full mesh interconexion network between all Border Nodes. If each island has two Border Nodes, primary and backup, cp-links can be established between primary Border Nodes and backup Border Nodes independently. If m is the number of islands, then $m(m-1)$ is the total number of cp-links (primary and backup). Possible scalability problems with control messaging, OSPF-TE routing information can be solved using OSPF area division in the MPLS network.

References

[1] CCAMP, http://www.ietf.org/html.charters/ccamp-charter.html

[2] K. Jumaki, “Interworking Requierements to Support operation of MPLS-TE over GMPLS Networks”, draft-ietf-ccamp-mpls-gmpls-interwork-reqts (work in progress), 2007.

[3] P. Pan, G. Swallow, A. Atlas, “Fast Reroute Extensions to RSVP-TE for LSP Tunnels”, RFC 4090, 2005.

[4] E. Mannie, “Generalized Multi-Protocol Label Switching (GMPLS) Architecture”, RFC 3945, 2004.

[5] G. Swallow, J. Drake, H. Ishimatsu, Y. Rekhter, “Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model”, RFC 4208, 2005.

[6] A. Farrel, J. P. Vasseur, A. Ayyangar, “A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering”, RFC 4726, 2006.

[7] K. Kompella, Y. Rekhter, “Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)”, RFC 4206, 2005.

[8] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, G. Swallow, “RSVP-TE: Extensions to RSVP for LSP Tunnels”, RFC 3209, 2001

[9] ITU-T G.8080, “Architecture for the automatically switched optical network (ASON)”, 2001.

[10] L. Berger, “Generalized Multi-protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions”, RFC 3473, 2003.

[11] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow. “RSVP-TE: Extensions to RSVP for LSP Tunnels”. RFC 3209 2001. Updated by RFCs 3936, 4420, 4874.

[12] K. Kompella, Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, 2003.

[13] S. Dharanikota, R. Jain, R. Hartani, D. Papadimitriou, Y. Xue, C. Brownmiller. “On Shared Risk Link Groups for diversity and risk assessment”, Contribution to T1 standards project, March 2001.