

Experimental validation of end-to-end GMPLS-enabled restoration in multi-domain transparent WSON

R. Martínez¹, R. Casellas¹, L. Velasco², F. Agraz², R. Muñoz¹, S. Spadaro²

¹Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Av. Carl Friedrich Gauss n7, 08860 Castelldefels, Spain

²Universitat Politècnica de Catalunya (UPC), Jordi Girona 1-3, 08034, Barcelona, Spain
{ricardo.martinez, ramon.casellas, raul.munoz}@cttc.es; {luis.velasco, agraz, spadaro}@tsc.upc.edu

Abstract: Two GMPLS restoration strategies operating at wavelength and link granularities are experimentally validated in a multi-domain WSON. The strategies are qualitatively compared in terms of resource usage, restorability and restoration time.

© 2009 Optical Society of America

OCIS codes: (060.4261) Networks, protection and restoration

1. Introduction

In Wavelength Switched Optical Networks (WSON), a link or node failure causes a huge amount of data loss. Fast and efficient recovery schemes are needed to minimize these effects and recover the disrupted services. This work concentrates on restoration schemes where the backup path is computed and set up after the failure is detected, localized and notified to the responsible node to recover the failed *lightpath*. This allows an efficient utilization of the network capacity at the expense, however, of compromising both the restoration success and time. In the last years, WSON restoration has been extensively studied. In general, these works focused on restoration within a single domain where the routing schemes for recovery are aware of the complete topology and network resource status, allowing high restoration efficiency. However, a complex WSON network may be formed by several domains (*multi-domain*) due to, for instance, scalability purposes. Hence, only abstracted topology and reachability information is shared among domains. While each domain is, in principle, responsible for the routing of the path segment traversing its respective network, the restoration of the end-to-end lightpaths may be far from the optimal yielding to an inefficient use of the overall network resources.

This work particularly addresses the multi-domain restoration in the GMPLS context. We consider that the lightpaths are *transparently* (i.e., within the optical layer) set up regardless of the traversed domains. That is, no optical-electronic conversions are used. In this regard, the working and backup lightpaths must fulfill the *wavelength continuity constraint* (WCC) since no all-optical wavelength converters (WCs) are placed.

In the GMPLS restoration, once a failure occurs, the selected nodes to repair the failed lightpath are notified with different information levels regarding the failed resources: no information, wavelength, and link node basis. This granularity allows devising different restoration strategies which operate at each of these information levels. The goal of this work is to validate the feasibility of the GMPLS protocols to restore lightpaths within multi-domain transparent WSON when using one of these restoration strategies. Specifically, when using failure information at either link or wavelength channel granularity. The validation is experimentally carried out within a multi-domain network connecting at the control plane level both CTTC ADRENALINE testbed® and UPC CARISMA testbed.

2. Problem description

We consider a multi-domain network formed by several OSPF-TE areas connected through the area 0. The nodes connecting two or more areas are the *Area Border Routers* (ABRs). The ABRs summarize the flooded topology information exchanged among domains. This leads a domain to have a limited visibility of the topology and resource status in the other domains. In consequence, the routing of lightpaths (*Label Switched Paths*, LSP in GMPLS) is attained in a per-domain basis. That is, for end-to-end LSPs spanning multiple domains, each ABR along the route computes, using its own intra-domain TE information, the segment of the LSP to the next egress ABR until the destination is reached. Therefore, a *route expansion* is required at each traversed ABR.

In the RSVP-TE signalling protocol, a set of Notify Request objects may be added to the Path and/or Resv messages indicating the node IP addresses to be notified when a LSP failure occurs. To this end, GMPLS uses the Notify message [1][2]. The node receiving such a message is termed as *Point of Repair* (PoR), and is one of the responsible to restore the failed LSP. We assume that the source and the ABRs traversed by a LSP act as PoRs.

In the Notify message, the inclusion of failure information (e.g., node, link, or wavelength) allows the PoR to avoid the failed resources when computing an alternate path [3]. Typically, such information is conveyed during the backup LSP setup as the *exclude Route Object* (XRO) [3] [4] within the Path message, so nodes doing a route expansion (e.g. ABRs) may avoid these failed resources. However, in multi-domain networks, it may be that a PoR within a given domain does not have sufficient topology information to compute a strict *Explicit Route* excluding the failed resources conveyed in the XRO object. This has an impact on the efficiency of the restoration and may require more advanced topology aggregation mechanisms using a hierarchical routing [5].

Let's use the Fig. 1 to analyze the impact of reporting the failure information to the PoRs when restoring a LSP. Two GMPLS domains (CTTC and UPC) are connected through the area 0. A working LSP is established along the path formed by nodes 4, 2, 8, 9 and 10. A failure occurs between nodes 9 and 10. Three different failure types may occur: a wavelength channel (e.g., problem on a receiver of any of the OXC ports), link (e.g., fiber cut) and node (i.e., optical switch). A Notify message is sent upstream by the node 9 reporting the failure to the PoR (i.e., node 8). This PoR is required to compute an alternate route, within the UPC domain, to detour the failure, maintaining the same wavelength. In the example, we assume that node 8 can not restore the LSP due to either the lack of resources or the WCC failure. Thus, the Notify message is sent to the upstream PoR of the route (i.e., node 2). Let's consider that the node 8 removes/filters from the Notify message the failure information to avoid sharing such information among domains. In consequence, node 2 is only aware that node 8 needs to be avoided as long as the same wavelength must be kept (due to the WCC), since it can not restore. It is worth noting that if wavelength converters were available in node 2, node 8 may be traversed with another wavelength. Thereby, node 2 is also unable to route the LSP within the area 0. Note that the other outgoing links from node 2 (i.e., with nodes 4 and 1) do not belong to the area 0. The Notify message is finally sent to the upstream PoR (i.e., node 4). In this case, node 4 (the source) can apply two routing policies: *pessimistic* and *optimistic*. In the former, the route detours the node 2 (i.e., path along the nodes 4, 3 and 1) and considers any available wavelength including the failed one. In the latter, the route does not detour node 2 but exclude the failed wavelength channel. Focusing on the pessimistic policy, a route expansion is needed at node 1. Since no XRO information is carried in the Path message, node 1 has two routing options: through node 8 or node 7. In the first, the restoration LSP would definitely fail. In the second option, the route expansion done at node 7 could route through either node 11 or node 9. Note that through node 9 the LSP would also fail. Hence, we observe that filtering information between domains may lead to block the restoration of LSPs even if a feasible backup LSP exists.

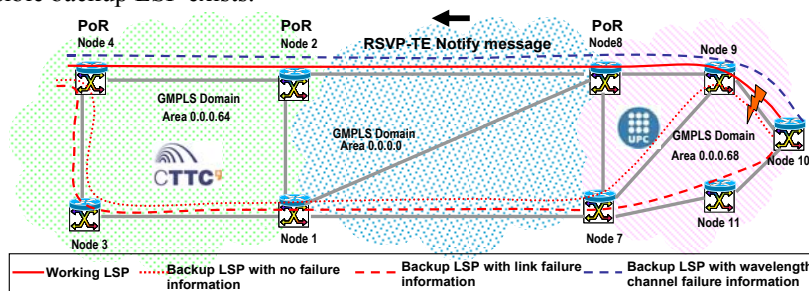


Fig. 1. Multi-domain LSP restoration with either no failure information and/or with link/wavelength-based information shared among domains.

In light of the above, to increase the LSP restorability, failure information needs to be exchanged beyond the domain boundaries. Accordingly, node 8 sends the Notify message to the neighbouring domain without filtering the failure information. In case of a link failure, the Notify message specifies that failed link (i.e., 9-10). Node 2 receives the Notify message and, again, node 2 can not route the LSP due to the WCC. Consequently, the Notify message is sent to node 4. Assuming the pessimistic policy, the backup path is computed along the nodes 4, 3 and 1 within the CTTC domain. The Path message contains the XRO to inform to the subsequent route expansions (nodes 1 and 7) about the nodes (i.e., 2 and 8) and the links (i.e., 9-10) to be avoided. This allows computing a feasible backup LSP along the nodes 4, 3, 1, 7, 11 and 10.

The above solution is adequate when the failures are either link or node. However, if the working wavelength channel fails in the link between nodes 9 and 10, the complete link should not be discarded for the backup LSP. Indeed, the backup LSP could be set up using a different wavelength channel on that link. The only restriction is that the WCC for the backup LSP must be satisfied. Thus, the source node may consider the optimistic approach (i.e., not detouring node 2 but without using the failed wavelength), based on a policy decision. In our example, the node 8 receives the Notify message including the link and the failed wavelength channel. Due to the WCC, node 8 sends the Notify message to the next PoR (node 2). Since node 2 is also unable to restore the LSP, a Notify message is finally sent to the node 4. As stated, node 4 may or may not reroute via 2. Let's assume not detouring is performed (optimistic policy). In this case, the Path message to set up the backup LSP includes the XRO which indicates the link and the failed wavelength channel. This information is taken into account within the UPC domain when performing the route expansion at node 8. In other words, until reaching such a node, the carried XRO information does not affect the needed path computations. However, at node 8, the wavelength channel failure information is used to constraint the path computation and the wavelength assignment algorithm. In the example, the computed segment path is formed by nodes 8, 9 and 10, in which the wavelength assignment algorithm excludes

the failed wavelength channel from the *Label Set* object. By doing this, usable wavelength channels on the link between nodes 9 and 10 are not discarded, and the LSP restoration may attain a better use of all the resources. In short, in a WSON with WCC, coarse granularity regarding link and node failures prevents the source node from deciding whether to reuse part of the failed path with an alternate wavelength or to exclude it completely. This choice is enabled by disseminating finer granularity failure information (i.e. at the wavelength level). Additionally, more efficient restoration mechanism can be deployed even in multi-domain scenarios knowing whether it was a failed wavelength/link/node without knowing exactly which one.

3. Experimental setup and validation

Fig. 2.a) shows the network used to validate the feasibility of the GMPLS protocols for dynamically restoring LSPs in a multi-domain WSON. Two domain networks at the control plane level, namely, the CTTC ADRENALINE and the UPC CARISMA testbeds are connected. The considered transport topology is shown in Fig. 1.

Besides verifying the intra- and inter-domain routing information, the tests validate the exchange of the RSVP-TE messages for notifying the same failure but at different information granularity: link and wavelength channel. Accordingly, two restoration strategies are used. The working LSP is set up through the nodes 4, 2, 8, 9 and 10 and occupies the label id. 637534212. The failure is generated at the working wavelength channel over the link between nodes 9 and 10. Fig. 2.b) depicts the RSVP-TE messages captured at node 8 when the failure information is notified at the wavelength channel basis. We observe that after the Notify message, the working LSP is torn down (see message frames 9-12). Next the backup LSP is set up along the same route as the working path (see message frame 13). In that Path message, the XRO indicates the resource at the wavelength channel (i.e., link and label id) to be avoided when performing the route expansion.

Table 1 gathers the results for both validated restoration strategies in terms of the computed backup LSPs along with their respective restoration time. As said, using failure information at wavelength channel granularity does not discard usable wavelength channels on specific links. Furthermore, it allows increasing the LSP restorability ratio as well as computing shorter (in terms of traversed hops and links) backup LSPs. The latter leads to both lower the restoration time and increase the likelihood of satisfying the WCC. However, these benefits are achieved at the expense of increasing the signaling control overhead and compromising the confidentiality among domains.

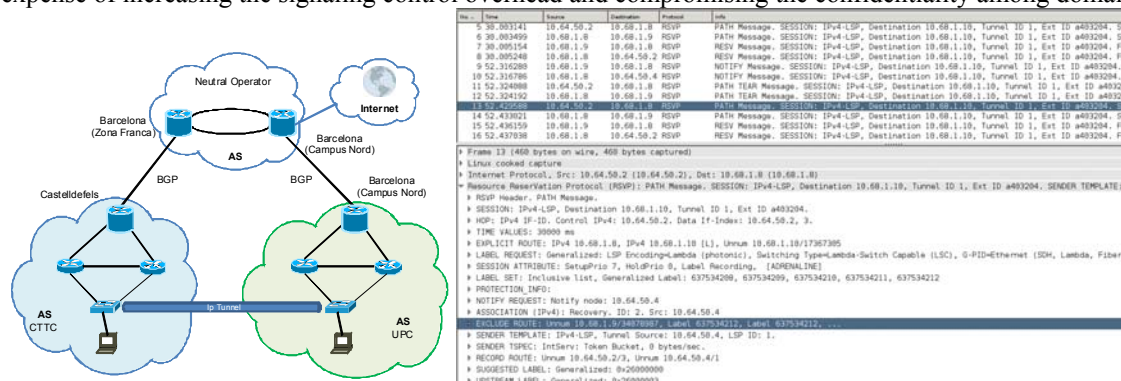


Fig. 2. (a) Simplified scheme of the IP network interconnection. (b) RSVP-TE messages at node 8 for the LSP restoration at wavelength channel.

Table 1. Numerical results

LSP Restoration using	Working LSP	Backup LSP	Restoration time (msec)
Failure information at wavelength channel basis	Nodes 4, 2, 8, 9 and 10; label id: 637534212	Nodes 4, 2, 8, 9 and 10; label id: 637534215	Around 135 msec
Failure information at link basis	Nodes 4, 2, 8, 9 and 10; label id: 637534212	Nodes 4, 3, 1, 7 11 and 10; label id: 637534215	Around 185 msec

5. Conclusions

Two restoration strategies (at either link or wavelength channel basis) for transparent multi-domain WSON are experimentally validated in a GMPLS-enabled control plane network constituted by two domains, namely, the UPC CARISMA and the CTTC ADRENALINE testbeds. Using the restoration at the wavelength channel level may lead to not only better address the WCC, but also to enhance the LSP restorability and reduce the restoration time.

6. References

- [1] J. P. Lang, *et al.* "RSVP-TE extensions in support of end-to-end GMPLS recovery", IETF RFC 4872 (2007).
- [2] L. Berger, *et al.* "GMPLS segment recovery" IETF RFC 4873, (2007).
- [3] A. Farrel, *et al.*, "Crankback signaling extensions for MPLS and GMPLS RSVP-TE", IETF RFC 4920 (2007).
- [4] CY. Lee, *et al.*, "Exclude routes - Extension to Reserouce ReserVation Protocol-Traffic Engineering (RSVP-TE)", IETF RFC 4874 (2007).
- [5] D. Papadimitriou, *et al.*, "Evaluation of existing routing protocols against ASON routing requirements", IETF RFC 4652 (2006).