First demonstration of SDN-controlled Multi-Layer Restoration and its advantage over Optical Restoration

Itay Maor(1), Ori Gerstel(1), Victor Lopez(2), Thomas Szyrkowiec(3), Achim Autenrieth(3), Bernd Pruessing(4), Nuno Borges(4), Fisher FU(5), Guiu Fabregas(6), Juan Pedro Fernandez -Palacios(2)

(1) Sedona Systems, 17 Hatidhar St, Raanana, Israel; orig@sedonasys.com
(2) Telefónica I+D/GCTO, Distrito Telefonica, c/Ronda de la Comunicación s/n, 28050 Madrid, Spain.
(3) ADVA Optical Networking SE, Fraunhoferstr. 9a, 82152 Martinsried, Germany
(4) Coriant GmbH, St.-Martin-Str. 76, D-81541 München, Germany
(5) Huawei Technologies Co.,Ltd., 518129, Shenzhen, China.
(6) Nokia, Maria Tubau 9, 28050 Madrid, Spain.

Abstract We demonstrate a full implementation of centrally orchestrated multi-layer restoration over commercial optical and IP gear. The process considers the behavior of the IP layer. Compared to optical restoration, the packet loss is 54% lower.

Introduction

Multi-layer restoration (MLR) has been known to enable very significant savings in a multi-layer IP and optical network [1,2]. However, this functionality has not yet seen commercial deployment due to lack of a full solution that covers network planning, automated restoration as well as a hitless reversion process. Moreover, since Service Provider networks are multi-vendor in nature, the solution must cover control interoperability across different optical and IP domains. Such interoperability has been attempted using distributed control plane mechanisms, such as GMPLS UNI [3], but it never provided sufficient functionality to enable such a full solution without a central view. The rise of SDN – in the form of centralized control over distributed systems, finally enables such a solution, as demonstrated in this paper.

In the past, most work on restoration involved simple optical restoration to circumvent the failure and restore lost IP capacity, and the specific limitations of the IP layer were ignored in the process. While [1] explains that the restoration process must involve the IP layer for failures that include router ports, it was not apparent until now that there is a need to include IP layer considerations in the restoration process even for pure optical failures (MLR-O in the terminology of [1]). This is the theoretical contribution of this paper. The experimental contribution is the demonstration and measurement of the first SDN based centrally controlled multi-layer restoration, to the best of the authors' knowledge. The only previous lab trial was done by Telefonica in [2], where GMPLS UNI was used. However, as previously mentioned, distributed GMPLS UNI architecture is insufficient to enable the functionality needed for commercial deployments.

What is multi-layer restoration and why is it needed

Multi-layer restoration for optical failures includes a multi-layer *planning phase*, in which a restoration plan is devised, and an *execution phase*, in which the plan is carried out. Multi-layer planning is a slow process of figuring out which connections should be restored when an optical failure occurs, what is their alternate route around the failure, and in what order should connections be restored. It takes place in advance before a failure occurs. The execution phase takes place after a failure has occurred. It must therefore be quick. This phase can be performed in distributed manner – autonomously by the optical layer nodes, or in centralized manner – by the multi-layer orchestrator. Both approaches have their pros and cons, which are outside the scope of this paper.

A key characteristic of optical switching is that it is slow – especially in large networks, where restoring a single connection can take tens of seconds. While this is likely not an issue for best effort low-priority traffic, high-priority traffic used for business services is expected to be restored by the IP layer, using existing fast mechanisms such as fast reroute [1]. Restoration of low priority traffic will also be attempted by the IP layer; however, such traffic may experience congestion until the optical layer switches the underlying connections away from the failure. The slow restoration of the optical layer also implies that the order of restoration matters. This important fact has been overlooked so far.

To minimize traffic loss it is important to restore

IP links that carry more traffic before restoring links that carry a lower traffic volume. It is also important to recognize that routers are often connected by parallel links that are bundled together – this is how IP networks scale to have multi 100Gbps connections between routers. The links in such a bundle should ideally be restored together, since the IP layer does not start using the bundle until a significant portion of its member links are up (subject to operator configuration). Both the magnitude of traffic carried over a link and its link bundle configuration implies that the order of restoring links matters. Not knowing this information, optical restoration might restore a portion of a bundle, then restore a portion of another bundle, so that none of these bundles becomes operational. It might also focus on links that carry a small amount of traffic, while neglecting links that carry a large amount of traffic. In both cases, the result is unnecessarily high traffic loss.

The multi-layer proof of concept

We demonstrate MLR in Telefonica's I+D/GCTO lab. It encompasses commercial routers and optical equipment of multiple vendors, both in the IP layer (Nokia 7750 SR, Juniper MX-240 and Huawei NE-40 routers) and the optical layer (Coriant hiT7300, ADVA FSP3000 & Huawei OSN 9800). The optical equipment is controlled through the vendor-supplied SDN controllers: Coriant Transcend™ SDN Transport Controller, ADVA Optical Virtualization Controller and Huawei Transport SDN Controller. The Nokia and Juniper routers are controlled directly via Netconf/YANG from the orchestrator, and the Huawei router is controlled via AGILE controller. The orchestrator that provides the overall solution is Sedona Systems Multi-Layer Application Platform (MAP), which interacts with the various controllers, collects topology and traffic information from each layer, puts together a multi-layer network view and abstracts the vendor equipment details for the use of a set of vendor-agnostic applications (see Fig. 1). One of these applications is a dedicated MLR planning application, while another one is a MLR execution application (in this case, centralized restoration was performed). The interface between MAP and the vendor controllers is a RESTful interface, typically closely aligned to the Transport API from the ONF, however the detailed model has been somewhat different between these interfaces.

Over this setup, we establish three connections between routers: a 10GE connection between the top router and the bottom right router, a 10GE connection between the bottom left router and the bottom right router, and a link bundle of two 1GE links between the bottom left router and the top router. We also establish three MPLS tunnels between the routers as shown in Fig. 2. The number in the figure reflects the reserved capacity of these tunnels. The setup also includes a Spirent traffic generator that injects traffic at the top router as shown in Fig. 1. The traffic is measured by polling SNMP counters every 5 seconds at the bottom left router.

We demonstrate in a multi-vendor and multi-layer environment:

a. Automatic L0-L3 network discovery and mapping.
b. Generation of a multi-layer restoration plan.
c. Review & simulation of the restoration plan.
d. Automated execution of the restoration plan.
e. Hitless revert after the failure has been fixed.

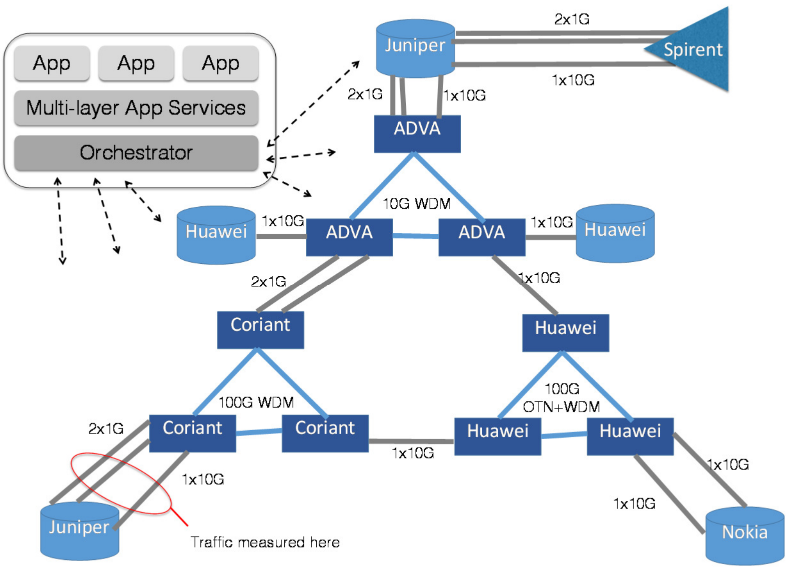

Fig. 1: Lab set-up

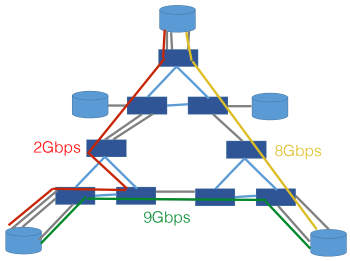

Fig. 2: IP flows before failure

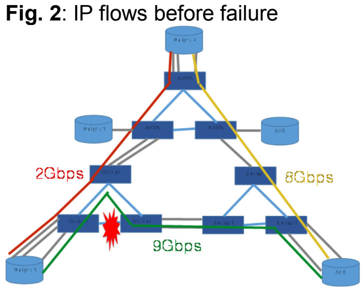

Fig. 3: IP flows after failure

The main scope of this paper is around capabilities b and d. In step b, the MLR Planning application computes a restoration plan for each possible failure. The plan defines which optical connections must be restored and in what order. The goal of this process is to find an order that maximizes the speed traffic is restored. To understand traffic behavior during this process, the application considers current traffic conditions and consults an IP simulation tool – in this case Cisco WAE Design. This planning process is an ongoing process, which updates the plan as traffic conditions and circumstances change.

The MLR Execution application is demonstrated by simulating a failure by taking down a DWDM link in the Coriant network as shown in Fig. 3. As a result, the Coriant Transcend™ SDN Transport Controller, notifies Sedona's MAP of the failure. This triggers execution of the restoration plan and MAP instructs the Coriant controller to take down the existing connections over the failed link and reestablish new ones around the failure. The order chosen for the process favors the 10GE link over the bundle of two 1GE links. This results in quick restoration of the 9Gbps MPLS tunnel, followed by the restoration of the 2Gbps tunnel. The traffic measurement results are shown in Fig. 4(a).

We also compare the effectiveness of MLR to pure optical restoration. To this end, we run a script that restores optical connections in an arbitrary order. In this case, it restores the two 1GE links before it restores the 10GE link. As a result, a small amount of traffic is restored first, while the bulk of the traffic must wait for the late restoration of the 10GE link. The measured traffic for this case is shown in Fig. 4(b).

It should be noted that in the lab optical switching is naturally too quick to display the slow restoration behavior that is experienced in the field, therefore we introduce an artificial delay before the restoration of each connection. In this case, the delay was set to two minutes. Such a delay is aligned with Telefonica's restoration time in field deployments. Note however that the absolute delay time is not fundamental to this experiment. What matters is the ratio between the traffic loss in both cases. In this case, the normalized traffic loss per second with MLR is approximately 19Gb +13Gb, versus a normalized traffic loss of 60Gb + 9Gb per second in the optical restoration case. The relative traffic loss for MLR is 32/69=46% of the traffic lost in the optical restoration scenario.

Summary

We demonstrated a realistic multi-layer and multi-vendor network, controlled by SDN controllers, and orchestrated by a multi-layer/multi-vendor SDN orchestrator. This setup was used to demonstrate multi-layer restoration and to compare its performance to that of pure optical restoration that is unaware of IP layer considerations. We measured the difference in traffic loss incurred during the restoration process. The traffic loss exhibited by multi-layer restoration was 54% lower than that of optical restoration, due to a selection of a connection restoration order that took into account the behavior of the IP layer.

Acknowledgements

This research was partially funded by the European Commission H2020 Research and Innovation program, ACINO project, Grant Number 645127, www.acino.eu.

References

[1] O Gerstel et al, "Multi-layer capacity planning for IP-optical networks," IEEE Comm. Magazine, 52 (1), 44-5, 2014.

[2] Carlos García Argos, "An Industrial Application of Multilayer Traffic Engineering Techniques." M.Sc. Thesis, 2011.

[3] RFC 4208, G. Swallow, J. Drake, H. Ishimatsu, Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", Oct. 2005.

(a) Recovery of traffic with MLR (b) Recovery of traffic with optical restoration

Fig. 4. Traffic measurements (in Gbps) with multi-layer restoration and without it