Policy-based Restoration in IP/Optical Transport Networks

M. Santuari[α], T. Szyrkowiec[β], M. Chamania[β], R. Doriguzzi-Corin[α], V. Lopez[γ], D. Siracusa[α]
[α]CREATE-NET, [β]ADVA Optical Networking SE, [γ]Telefónica I+D

I. INTRODUCTION

Restoration in transport networks is typically facilitated using reactive techniques at different layers, namely optical and IP restoration [1]. Optical restoration involves re-routing an existing optical connection (i.e., a lightpath) around a failure (e.g. link, amplifier, switch and transponder failures) in the optical layer. This strategy is efficient in terms of resource utilization, as backup resources are reserved dynamically after the failure and therefore are not blocked during normal operation. However, equipment reconfiguration and power equalization processes in the optical domain are relatively slow (order of seconds), and are thus not suitable for time critical services.

IP restoration also re-computes path around failures (e.g. IP links, router ports) and re-uses existing IP links to support the restored traffic. Reconfiguration in IP is fast (less than 50 milliseconds), but requires that alternate paths have sufficient capacity to support the re-routed traffic, which in turn may imply blocking of resources in the optical domain, therefore requiring additional active lightpaths.

The two restoration strategies present trade-offs in the form of cost (additional lighpath provisioning vs. re-use of existing lighpaths), responsiveness (seconds vs. milliseconds) and offered capacity (same as primary lightpath vs. shared spare capacity on the backup lightpath), and can complement each other to cater to the needs of emerging high-bandwidth services related to 5G and cloud applications. However, in order to be effective, network intelligence is required to identify the characteristics of a failure and the affected services, and orchestrate restoration on the best-suited (IP/optical) layer.

This paper[1] proposes the first demonstration of an IP/Optical SDN control solution for transport networks, called network orchestrator, which orchestrates IP or optical restoration based on the policy explicitly requested by the client application. The policy is communicated via intents, as part of the constraints that must be satisfied for a service. The orchestrator uses these intents to identify the restoration mechanism to be employed in case of a failure. The proposed orchestration solution allows the operator to effectively utilize network resources while ensuring that the survivability constraints of requested services are met.

II. ARCHITECTURE AND TESTBED OVERVIEW

The proposed network orchestrator is built on top of the ONOS framework [2] and the core components of the orchestrator are presented in Fig. 1. The orchestrator exposes a North-Bound Interface (NBI) towards applications or Network Management Systems (NMSs), giving them the possibility to request network services with specific requirements (in this case, the desired restoration strategy). Requirements, expressed by means of intents, are then translated into service configurations by the intent framework, which relies on the knowledge of the IP/optical multi-layer topology. The South-Bound Interface (SBI) consists of the Providers and Protocols, as well as the technology-specific Drivers (according to the ONOS architecture), which extract information from the underlying infrastructure to build the network topology. The SBI is also responsible for pulling state changes and alarm notifications from network and pushing configurations computed by the orchestrator onto the network infrastructure.

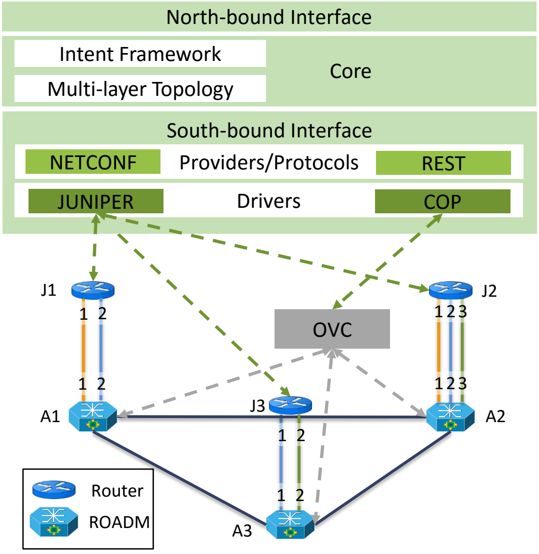

Fig. 1: High level orchestrator architecture.

For the proposed demonstration, the ONOS intent compiler has been modified to accommodate requests for IP or optical restoration. A new ONOS driver has been implemented to interface directly with the Juniper routers. The driver facilitates the discovery of devices and supports alarms notification to identify operational state changes in the IP network infrastructure. Configuration of IP routers is facilitated over the NETCONF protocol. With respect to the optical domain, the orchestrator interacts with the ADVA optical controller, called OVC (Optical Virtualization Controller). New ONOS driver and providers have been implemented to interact with the OVC in order to discover devices and links, provide alarm notifications and request the configuration of optical lightpaths. The orchestrator uses the Control Orchestration Protocol (COP) [3][4] to communicate with the OVC, which exposes a network model that is a subset of the one defined in the ONF T-API [5] and supports provisioning of lightpaths in the optical domain. A custom REST extension has also been developed to provide alarm notifications from the optical infrastructure to the orchestrator.

---

[1] This work is supported by the EU H2020 ACINO project, Grant Number 645127, www.acino.eu.

The demonstration will run on top of a physical IP/optical testbed located in Telefónica premises in Madrid. The testbed setup, as presented in Fig. 1, consists of 3 Juniper MX240 routers (J1, J2, J3) connected via multiple 1GBE and 10GBE interfaces to 3 ADVA FSP3000 R7 WDM nodes (A1, A2, A3). The WDM nodes have the same configuration with directionless add/drop capabilities and are interconnected with each other in a ring configuration. In Fig. 1, IP and optical ports are numbered, while the color code (orange, blue, green) identifies distinct wavelengths that are used on these interfaces. Therefore, a lightpath can be established between J1/1 and J2/1 (orange), J2/3 and J3/2 (green), and any pair of interfaces J1/2, J2/2 or J3/1 (blue), where Jx/y represents router Jx port y. The testbed setup provides diversity with respect to optical paths and capacity between any pair of routers.

III. DEMONSTRATION

In the demonstration we assume that an interconnection between an IP router and a ROADM fails, generating a port down alarm sent by the OVC and the IP router. According to the policy specified on the installed intent, the orchestrator attempts to restore the service at the optical or the IP layer. The combined sequence diagrams are shown in Fig. 2, the operations for optical and IP restorations are described below.

A. Optical Restoration

An intent requesting a service with a high bandwidth demand might require the setup of a dedicated lightpath to serve the traffic in case the existing connection fails. The intent for the primary connection (Intent 1) is created between J1 and J2, using interfaces J1/1 and J2/1. Please note that the installation of Intent 1 is not shown in the workflow presented in Fig. 2 due to space constraints, but it will be part of the demonstration. After the successful installation of the intent, a failure is triggered on the link connecting J1/1 and A1/1, and *Port_Down* notifications are received by the orchestrator from the OVC and J1. Based on these notifications, the orchestrator identifies all affected intents, and recompiles the intent for the failed services (only Intent 1, in this case). The intent compiler identifies an alternate optical path that is available between J1 and J2 using interfaces J1/2 and J2/2, and initializes the setup of the backup connection. The orchestrator triggers the creation of new lightpath between the optical ports connected to the routers, which is translated into configuration commands in the optical domain. As soon as the lightpath setup is completed, the orchestrator initializes the configuration of the new routes on J1 and J2. The successful setup of routes finalizes the optical restoration process.

B. IP Restoration

Operators may require to reroute best-effort services over existing IP links in case of failure. In this second scenario, a full mesh of IP links is established between J1, J2 and J3: Intent 1 (J1/1 - J2/1), Intent 2 (J1/2 - J3/1), Intent 3 (J3/2 - J2/3). In particular, Intent 1 selects IP restoration policy (the policy for the other intents is not relevant for the demonstration). The initial installation of intents is not shown in Fig. 2, but it will be part of the demonstration. After the successful installation of the intents, a failure is again triggered on the link connecting J1/1 and A1/1, leading to *Port_Down* notifications. These notifications are received by the orchestrator, causing the failure of the IP link between J1 and J2. Based on the policy defined in the failed intent (Intent 1), the intent compiler is constrained to re-use existing IP links to serve the traffic. As a result, restoration is performed by configuring static routes on the IP routers in order to re-route the traffic from J1 to J2 via J3 and vice versa. However, if either Intent 2 or Intent 3 are not installed (i.e. there is no connectivity at the IP layer after the failure), the Intent 1 recompilation fails, as the setup of a new optical link is not supported for this best-effort service.

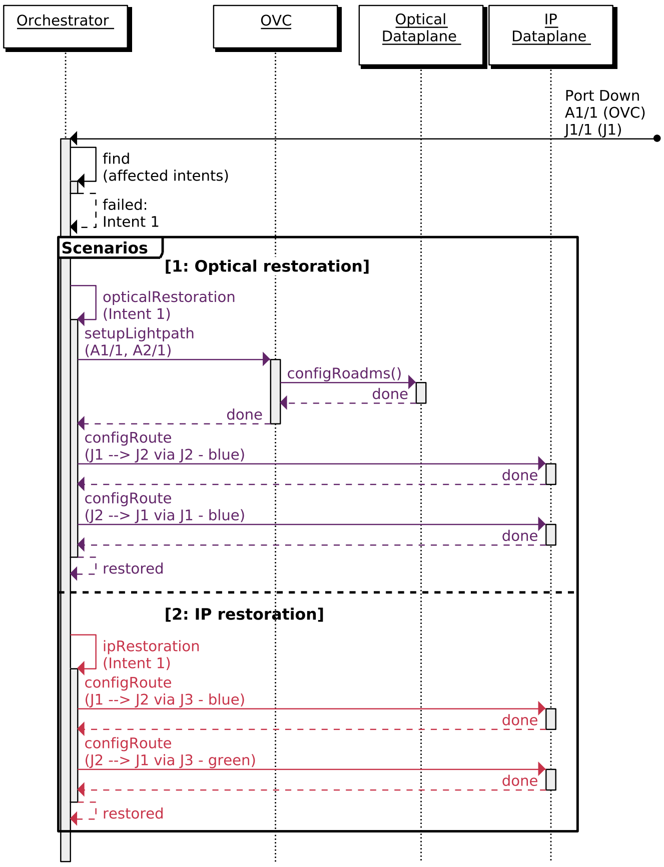

Fig. 2: Combined sequence diagrams for the optical and IP restoration scenarios.

IV. CONCLUSION

We presented a network orchestrator for multi-technology IP/optical transport networks and demonstrated its capability to apply policies defined in service requests to dynamically perform restoration at the IP or optical layer after a failure occurs. The capability to dynamically evaluate and perform restoration on a per-service level is essential for the effective utilization of network resources in the next-generation transport networks.

REFERENCES

[1] O. Gerstel et al., “Multi-layer capacity planning for IP-optical networks,” *IEEE Communications Magazine*, vol. 52, no. 1, pp. 44–51, 2014.

[2] P. Berde et al., “ONOS: Towards an Open, Distributed SDN OS,” in *HotSDN*, 2014.

[3] R. Vilalta et al., “The need for a Control Orchestration Protocol in research projects on optical networking,” in *Networks and Communications (EuCNC), 2015 European Conference on*, 2015, pp. 340–344.

[4] “Control Orchestration Protocol,” https://github.com/ict-strauss/COP.

[5] ONF, “Functional Requirements for Transport API,” v16, 2015.