Data Center Disaster Recovery Through an Open ROADM SDN-enabled Network

Dhruv Bhardwaj[1], Shweta Vachhani[1], Martin Birk[1], Christophe Betoule[2], Gilles Thouenon[2], Ahmed Triki[2], Guillaume Lambert[2], Olivier Renais[2], Marc Lyonnais[3], Joe Kunz[3], Tommy Locke[3], Greg Sutherland[4], YongChao Fan[4], Jian Zhang[4], Mark Dalglish[4], Lynne Hogan[4], Behzad Mirkhanzadeh[5], Ali Shakeri[5], Gi Vania[5], Miguel Razo[5], Marco Tacca[5], Tim Culver[5], and Andrea Fumagalli[5]

AT&T[1], Orange[2], Ciena[3], Fujitsu[4], The University of Texas at Dallas[5]

March 2019

OFC
The future of optical networking and communications is here.

Introduction

Data centers are critical components for the correct and efficient functioning of a number of business sectors. Any interruption in a data center operation can virtually bring down an entire business to its knees, unless an efficient backup strategy is devised. The use case envisioned in this demonstration consists of a backup data center that can temporarily substitute the primary data center in the presence of an imminent disaster. The backup and primary data centers, which must be located at a sufficient geographical distance from one another to guarantee spatial diversity, are interconnected by a *programmable optical fiber network consisting of multi-vendor Open ROADM equipment*. Upon detection of an imminent disaster affecting the primary data center, the *PROnet SDN Orchestrator* executes an automatic procedure whose goal is to live migrate the Virtual Machines (VMs), currently running at the primary data center, to the backup data center in the shortest possible time interval. To achieve this goal, the PROnet Orchestrator makes concurrent use of three open source platforms: *Open ROADM*, *OpenFlow*, and *OpenStack*. The first platform is used to dynamically establish a high data rate optical circuit between the two data centers, the second platform is used to establish a data flow between the top of the rack switches of two compute node racks, one at the primary and one at the backup data center, and the third platform is used to execute the live migration of the VMs over said data flow and optical circuit. The procedure terminates once all of the VMs are successfully running at the backup data center and both the data flow and optical circuits are relinquished.

Open ROADM MSA

The goals of the Open ROADM Multi-Source Agreement (MSA) are 1) the disaggregation and opening up of traditionally proprietary ROADM systems, and 2) the SDN-enablement of traditionally stationary ROADMs.

In this demo Ciena and Fujitsu's ROADMs and transponders are interchangeably interconnected to demonstrate Single-Wavelength and Multi-Wavelength interoperability.

Optical Specifications

Based on the optical functions, two optical specifications can be defined, 1) ***Single-Wavelength*** (or W) to define how **Pluggable Optics** or **Transponders** interoperate and 2) ***Multi-Wavelength*** (or MW) to define how **ROADMs** interoperate.

ROADM Systems Disaggregation

ROADM system disaggregation is accomplished by both ***hardware disaggregation*** and ***functional disaggregation***. The former refers to defining a common set of hardware provided by multiple vendors. The latter refers to defining optical functions. This demonstration focuses on functional disaggregation only, specifically targeting three optical functions defined by the Open ROADM MSA: 1) ***ROADM***, 2) ***Transponder***, and 3) ***Pluggable Optics***.

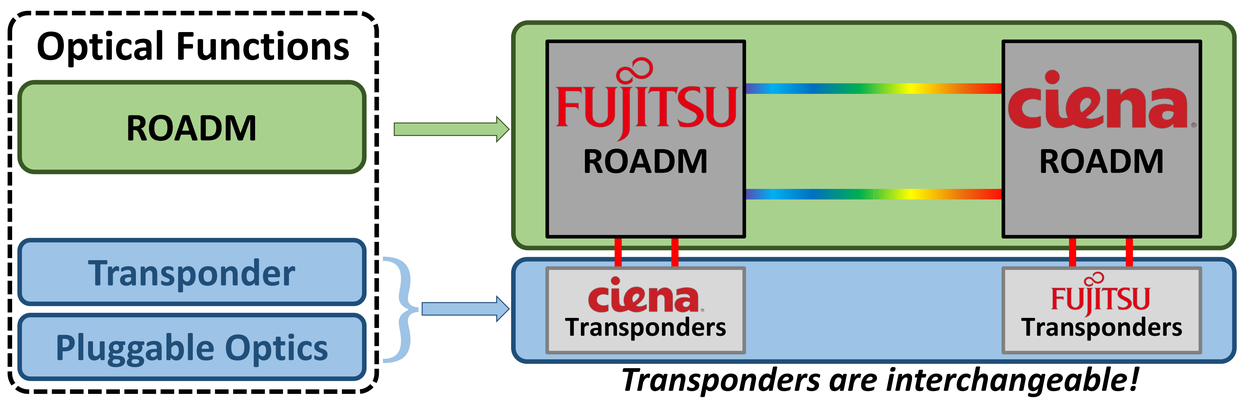

Integration of Open ROADM TransportPCE and PROnet SDN Orchestrator

All of the functions implemented in this demonstration are controllable through an open standards-based API written in the data modeling language YANG. This API is accessed through an SDN controller using NETCONF interface for service provisioning and network monitoring. Two ROADMs are controlled by the ***Open ROADM TransportPCE*** plugin. This is an open source plugin implemented in the OpenDaylight (ODL) controller. The required disaster recovery steps are sequentially launched by the ***PROnet SDN Orchestrator***, which first provisions a wavelength circuit (at 100 Gbps) between the primary data center, affected by the imminent disaster, and the backup data center, then creates end-to-end Ethernet flows between the servers in the two data centers, requests the live migration of the VMs from the primary data center to the backup data center through OpenStack, and finally relinquishes the network resources, when they are not any longer used.

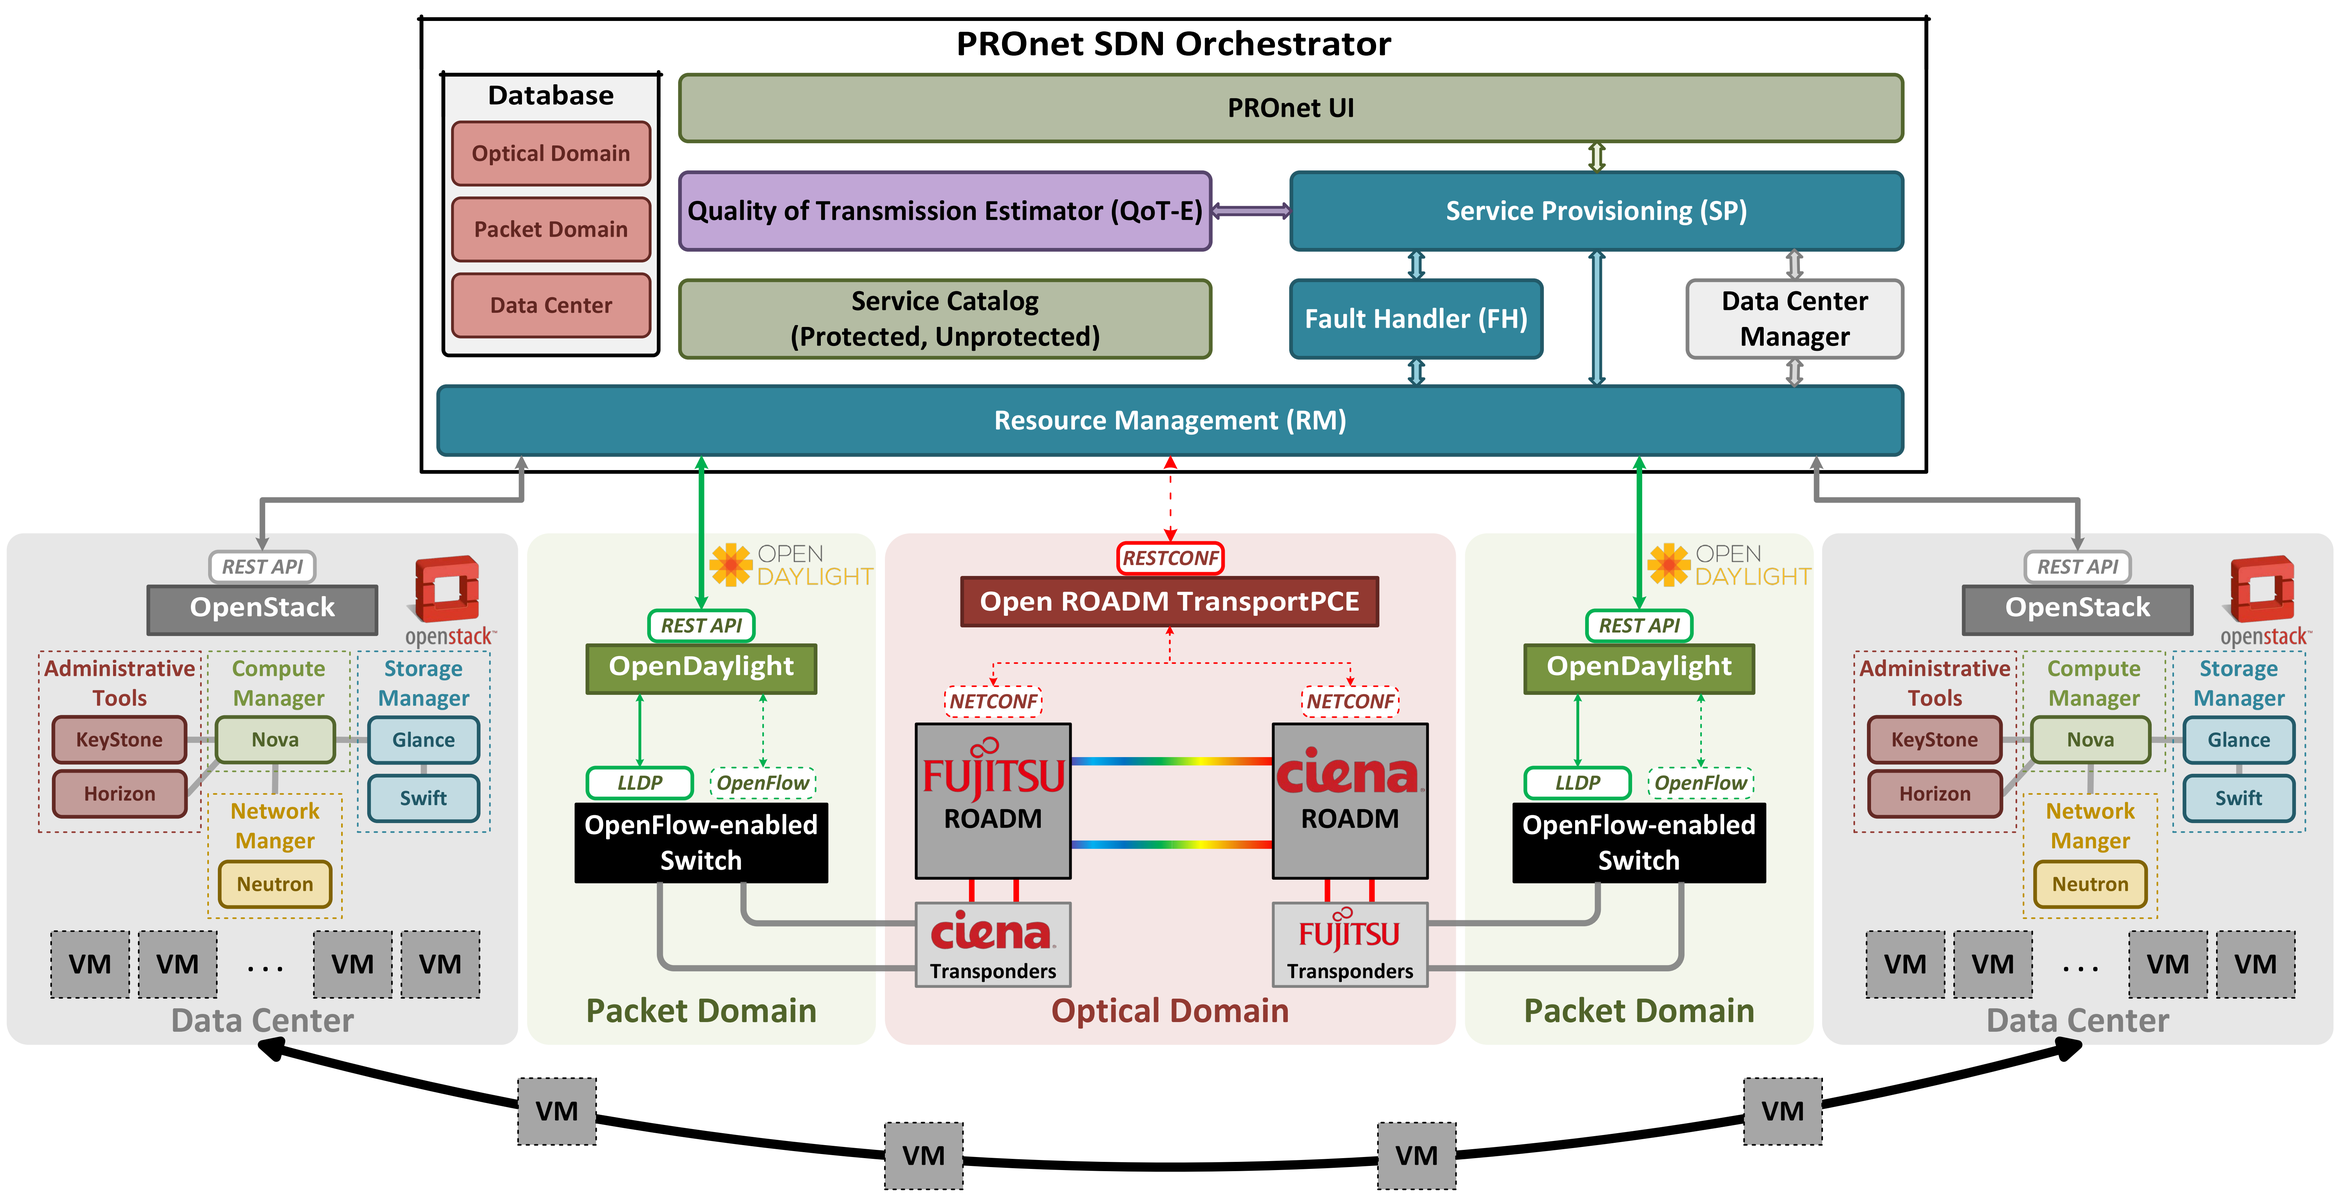

This work is supported by the Erik Jonsson School of Engineering and Computer Science, Texas Photonics Center, and NSF under grants CNS-1405405, CNS-1409849, ACI-1541461, CNS-1531039.

Some of the descriptions used in this poster are obtained from *http://www.openroadm.org/download.html*