

Interoperability of Innovative Capabilities in OpenROADM SDN-enabled Network

Martin Birk¹, Joe Kunz², Gabriele Galimberti³, Robert Sparks⁴, Greg Sutherland⁵, Davide Forin⁶, Mark Wiltz⁷, Olivier Renais⁸, Behzad Mirkhanzadeh⁹, Tianliang Zhang⁹, Joey White-Swift⁹, Gi Vania⁹, Miguel Razo⁹, Marco Tacca⁹, and Andrea Fumagalli⁹
 AT&T¹, Ciena², Cisco³, ECI⁴, Fujitsu⁵, Infinera⁶, Juniper⁷, Orange⁸, The University of Texas at Dallas⁹

March 2020



Introduction

Programmable optical fiber networks – among other things – are best suited for supporting bandwidth-greedy applications that do not need 24/7 connectivity. In this demonstration we consider two such applications: 1) massive and parallel live migrations of Virtual Machines (VMs) from a primary data center to a backup data center and 2) distributed machine learning training tasks. To support these applications the programmable optical fiber network is dynamically reconfigured to provide high-speed (100Gbps or 10Gbps) and dedicated optical circuits across a metropolitan area to facilitate the necessary data transfer that takes place between the involved end client nodes for just the duration time of each application. The programmable optical fiber network consists of interoperable multi-vendor OpenROADM equipment (Reconfigurable Optical Add-Drop Multiplexer (ROADM) nodes, transponders, and switchponders) that is provided by six equipment vendors. This multi-vendor optical network is controlled by an open source SDN (Software Defined Network) controller called TransportPCE – available as part of the OpenDaylight distribution – that natively supports OpenROADM-compliant network elements through the APIs defined in and published by the OpenROADM MSA. Through its northbound RESTCONF interface, the TransportPCE controller communicates with the PRONet (Programmable Optical Network) SDN orchestrator which provides a single point of control and coordination of resources across three domains, namely the optical fiber layer, Ethernet layer, and datacenter compute nodes. The PRONet orchestrator automatically and sequentially initiates actions that are required to execute datacenter backups, massive VM (Virtual Machine) live migrations, and other applications that require the high-rate and low-latency data transfer provided by OpenROADM equipment. While the OpenROADM network products that are integrated into this demo are already in use by network operators in their live networks, this is by far the largest public interoperability demonstration of the key and innovative capabilities that are enabled by the OpenROADM MSA.

OpenROADM MSA

The goals of the OpenROADM 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.

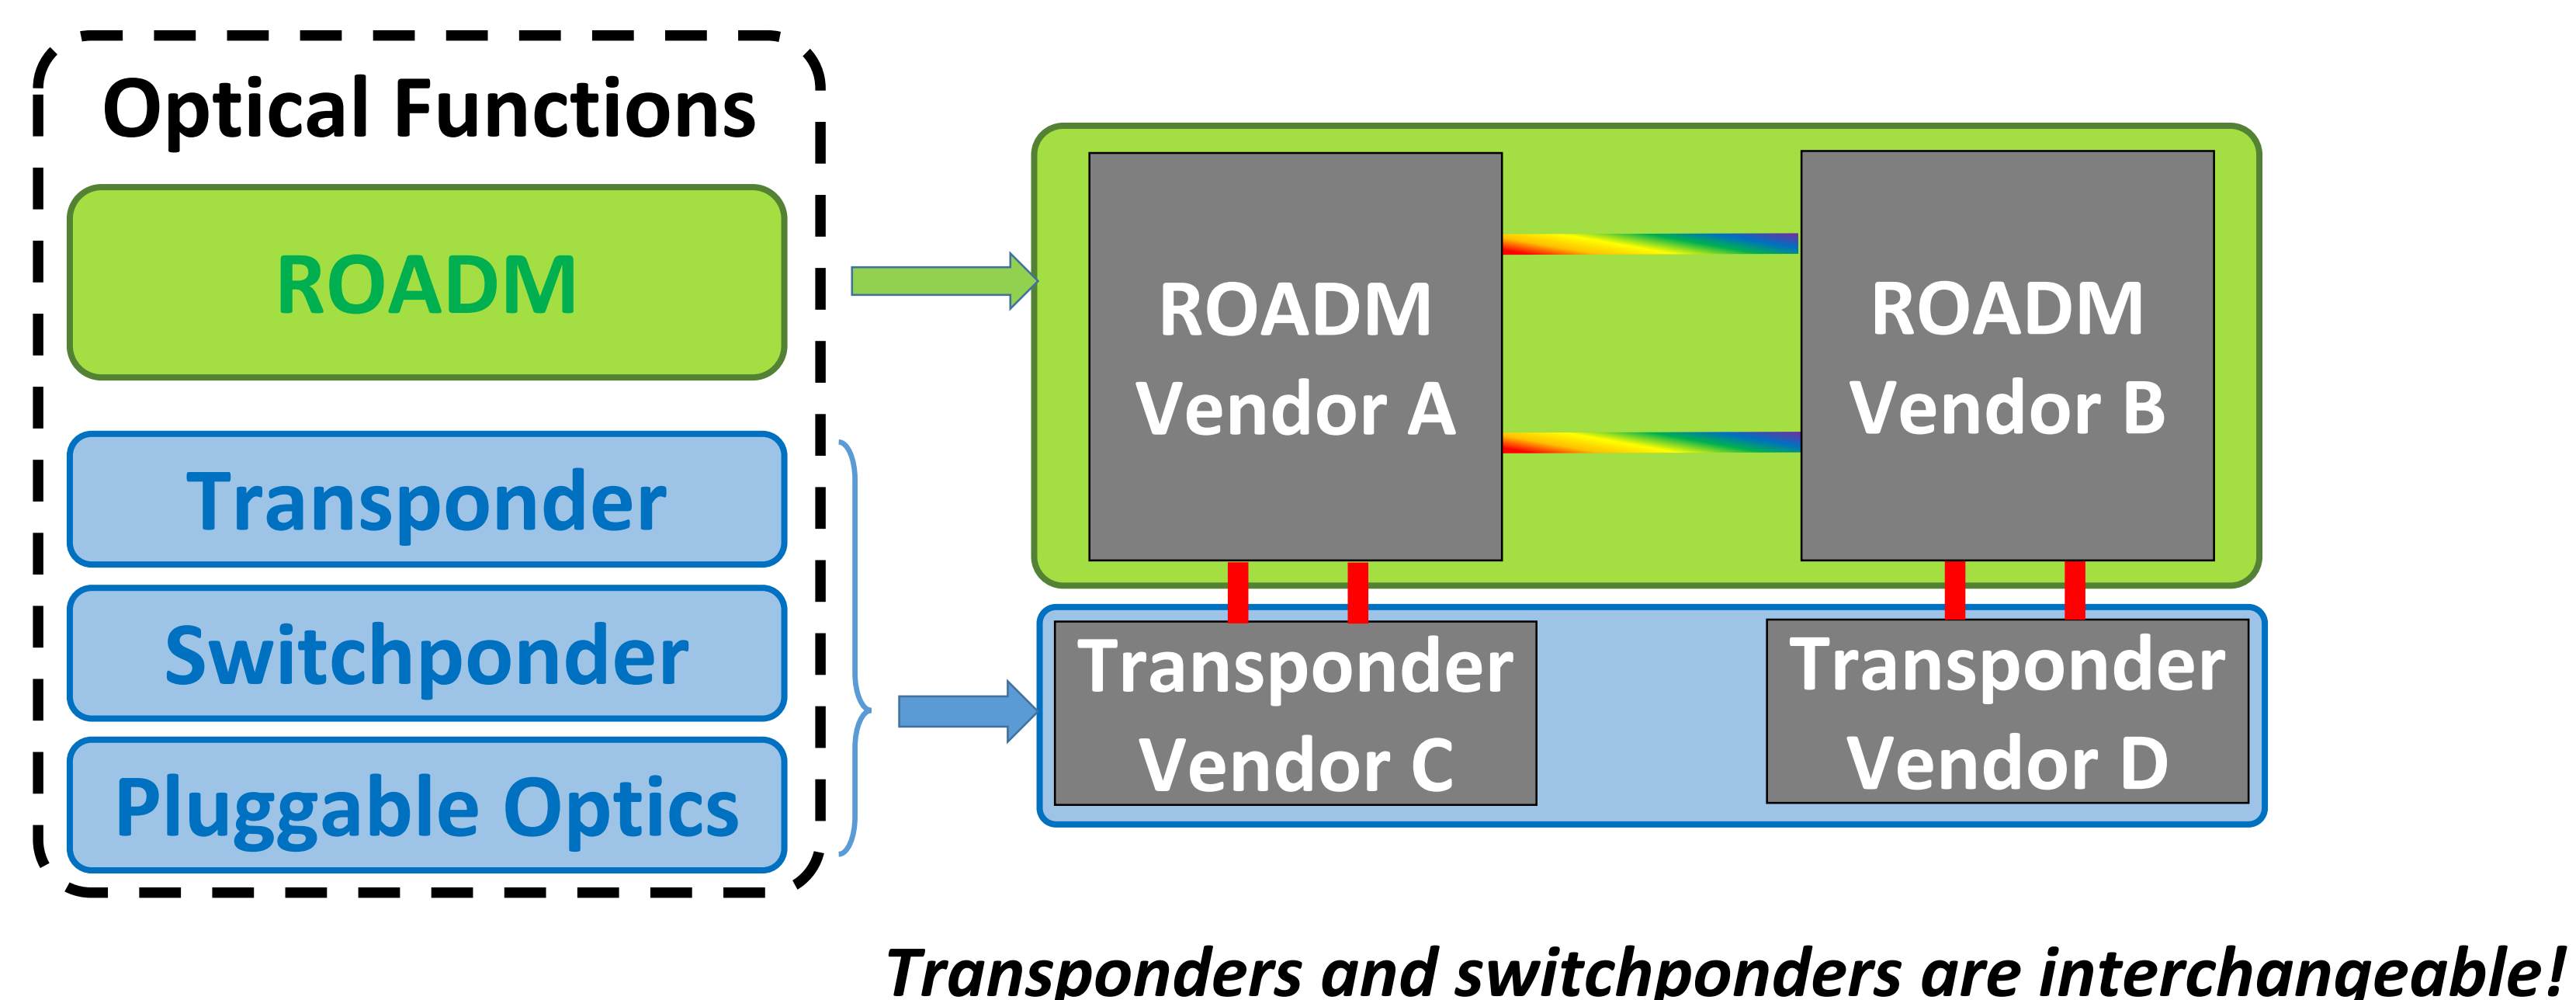
In this demo four Ciena and Fujitsu's ROADM nodes along with a number of transponders and switchponders (provided by all six participating vendors) are interchangeably interconnected to demonstrate Single-Wavelength and Multi-Wavelength interoperability.

Optical Specifications

Based on the optical functions, two optical specifications can be defined, 1) **Single-Wavelength** (or W) to define how **Pluggable Optics**, **Transponders** and **Switchponders** 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 four optical functions defined by the OpenROADM MSA: 1) **ROADM**, 2) **Transponder**, 3) **Switchponder**, and 4) **Pluggable Optics**.



Integration of OpenROADM TransportPCE and PRONet SDN Orchestrator

All of the optical 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. Four ROADM nodes, sixteen 100Gbps transponders, and five switchponders (line rate of 100Gbps and client rates of 1Gbps and 10Gbps) are controlled by the **OpenROADM TransportPCE** plugin. This is an open source plugin implemented in the OpenDaylight (ODL) controller. Six Juniper QFX OpenFlow-enabled Ethernet switches are controlled by an OpenDaylight controller. A total of forty repurposed Stampede compute nodes – assembled to form four compute sites – are controlled through OpenStack. The **PRONet SDN Orchestrator** – interfaced with the TransportPCE, OpenFlow ODL controller, and OpenStack – automatically executes the VM live migration procedure by sequentially first provisioning a wavelength circuit (at 100Gbps) between the selected primary data center – affected by the imminent disaster – and the selected backup data center, then creating end-to-end Ethernet flows between the servers in the two data centers, requesting the live migration of the VMs from the primary data center to the backup data center, and finally relinquishing the network resources (optical circuit and Ethernet flows), when they are not any longer used. The Orchestrator also coordinates with Kubernetes to deploy distributed machine learning training tasks across sites.

