Global Transport SDN Prototype Demonstration

OIF/ONF White Paper

October 7, 2014

Table of Contents

EXECUTIVE SUMMARY:

The interest in software-defined networking (SDN) will translate into a global market worth several billion dollars within the next few years. The increasing need for efficient infrastructure and mobility, as well as the growing popularity of cloud services, will drive this growth.

There are three areas motivating the main markets for SDN: enterprises, cloud services providers, and telecommunications services providers. However, cloud service providers are expected to be the fastest growing market segment throughout the years. SDN's ability to reduce OPEX and CAPEX while enabling the delivery of new services/revenue will spearhead its use by cloud service providers.

SDN and virtualization will shape the future of telecommunications by improving the network's ability to dynamically adapt to the needs of applications and services. The OIF and ONF seek to assure a seamless evolution to interoperable transport networks and a seamless operation of heterogeneous networks in an SDN architecture.

Participating OIF and ONF member carriers and system vendors joined forces to test prototype transport SDN technology in real-world scenarios. The Global Transport SDN Prototype Demonstration was supported by nine system vendors jointly conducting tests with five carriers in their labs across Asia, Europe and North America. OpenFlow™ extensions developed in the ONF Optical Transport Working Group were prototyped and tested in the demo in both CDPI and CVNI forms. Vendors also tested prototypes of Controller Northbound interfaces for Service Request and Topology functions in development by the OIF. The framework of the demo is cloud-bursting, or application-based bandwidth-on-demand, between data center sites. Testing based on this real-world use case illustrates potential deployment of Transport SDN technology, common interfaces required, needs for interoperability and any operational challenges.

The experiences from the demo will be shared across the industry to help develop Implementation Agreements and specifications. The ultimate goal is to accelerate the deployment of practical, programmable transport networks that enable a new era of dynamic services.

Introduction

Wide area network traffic, now acknowledged to be dominated by video flows, is growing unabated. This is exacerbated by dynamic and shifting traffic patterns of mobile and cloud-based services. Today's static, manual optical transport networks are not optimized to meet customer needs for flexible, on-demand services.

This situation challenges transport network operators and the supplier ecosystem to:

- Improve the network's ability to dynamically adapt to the needs of applications and services (residing mostly in data centers); i.e. make the network more programmable;
- Increase network efficiency and agility i.e. improve QoS/QoE, decrease OpEx, maximize revenue generation, improve ROI;
- Transform today's static networks while leveraging installed base/investments.

Many carriers are in the process of moving from SDN/NFV proof-of-concept projects to working with vendors in the development and 'productization' of software that will become the basis for commercial deployments.

Software-defined Networking (SDN) and virtualization promise to simplify optical transport network control by adding management flexibility and programmatic network element control to enable the rapid services development and provisioning. The centralized SDN-based network-wide management and control can drive improvements in network efficiency and speed in terms of service acceleration. SDN may also reduce the total cost of ownership (TCO) for optical switches by moving control and management planes from embedded processors to general-purpose common off-the-shelf (COTS) hardware and virtualized software.

To address this opportunity the Optical Internetworking Forum (OIF) and the Open Networking Foundation (ONF) joined forces to test prototype transport SDN technologies with the shared goal to make transport networks more programmable in order to enable a new era of dynamic services.

To this end, the OIF/ONF Global Transport SDN Interoperability Demonstration testing was conducted over several weeks in a number of global carrier labs, leveraging the OIF's carrier members representation, knowledge of transport networks, and worldwide interoperability testing experience for optical equipment, combined with ONF's leadership role for the OpenFlow™ protocol and SDN architecture. OpenFlow extensions for optical transport developed in the ONF Optical Transport Working Group and network APIs under development by the OIF were prototyped in the demo.

Demo highlights:

- Application: Cloud bursting, aka bandwidth on-demand, over optical networks characterized by:
    - Short-lived, high-volume traffic between data centers
    - Implemented using OTN connections created via a central interface
    - Utilizing Ethernet service over OTN as the data plane
- Protocol Features:

  - Subset of ONF Optical Transport Working Group (OTWG) Extensions
  - CDPI and CVNI interoperability
  - Experimental encoding of Extensions
- Controller NBIs (OIF lead):
  - Multi-domain (East/West or Controller hierarchy)
  - Network application interfaces (APIs) (Service Request, Topology Request)

This joint work is an initial step in addressing and offering solutions to the practical issues of implementing transport SDN in commercial networks. The OIF is currently working on several initiatives supporting Transport SDN including a carrier-driven Requirements Document and an SDN Framework Document identifying SDN application programming interfaces for a carrier environment. ONF currently has a number of carrier-focused initiatives underway, including SDN transport, mobile and wireless network applications, carrier-grade SDN, and large-scale network migration.

Demonstration Set-up

Testing Overview

This test addresses a complete SDN architecture as might be deployed by a carrier. The SDN architecture is shown in Figure 1 below[1]:

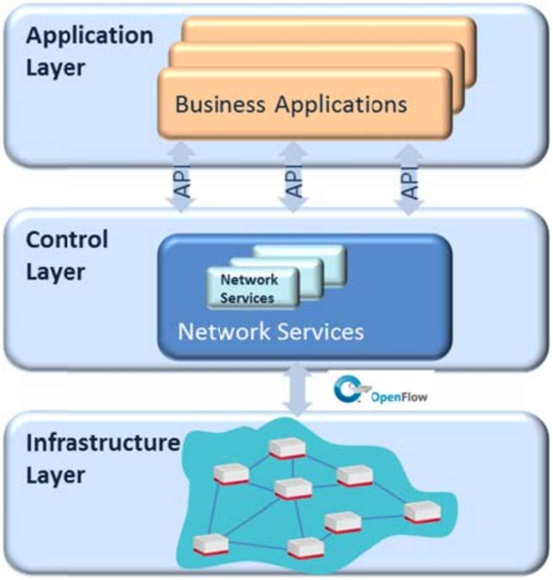

Figure 1: SDN Architecture

---

[1] https://www.opennetworking.org/images/stories/downloads/sdn-resources/solution-briefs/sb-of-enabled-transport-sdn.pdf

SDN identifies interfaces separating data plane and control plane, and control plane from application. The interface between data plane and control plane is termed the Southbound Interface or SBI, the interface between control and application is termed the Northbound Interface or NBI. The demonstration prototyped all three layers in the SDN architecture, as well as standard SBI and NBI interfaces between layers.

Implementation of the SDN model layers was as shown in Figure 2:

- The Infrastructure Layer consisted of multiple domains of Network Elements, including Ethernet switches, OTN ODU switches and OTN OCh switches;
- The Control Layer consisted of the Domain Controllers and in some cases a Parent or Network Operator Controller, using OpenFlow or other protocols as the SBI;
- The Application Layer consisted of the Network Orchestrator, an application that uses the NBI offered by the Control Layer to obtain topology information and request service across the network, either by a customer or by an internal network process.

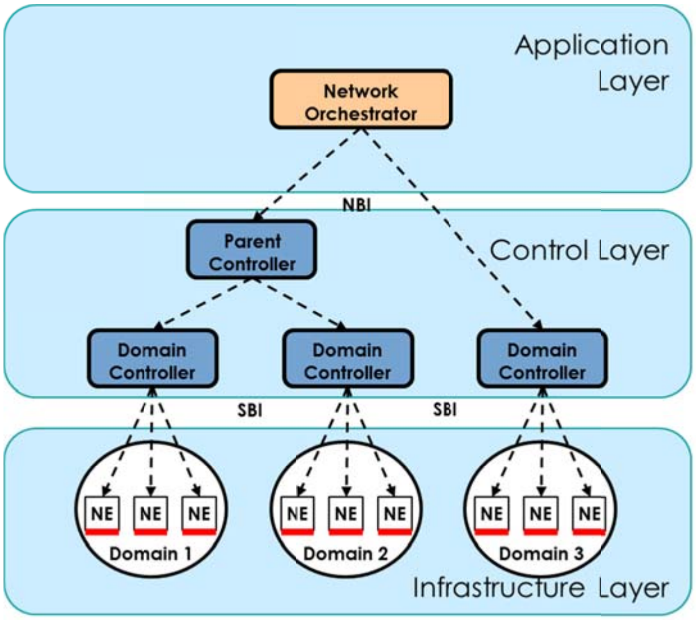

Figure 2: SDN Layer Implementation

Worldwide Test Topology

The Worldwide Test Topology is shown below in Figure 3:

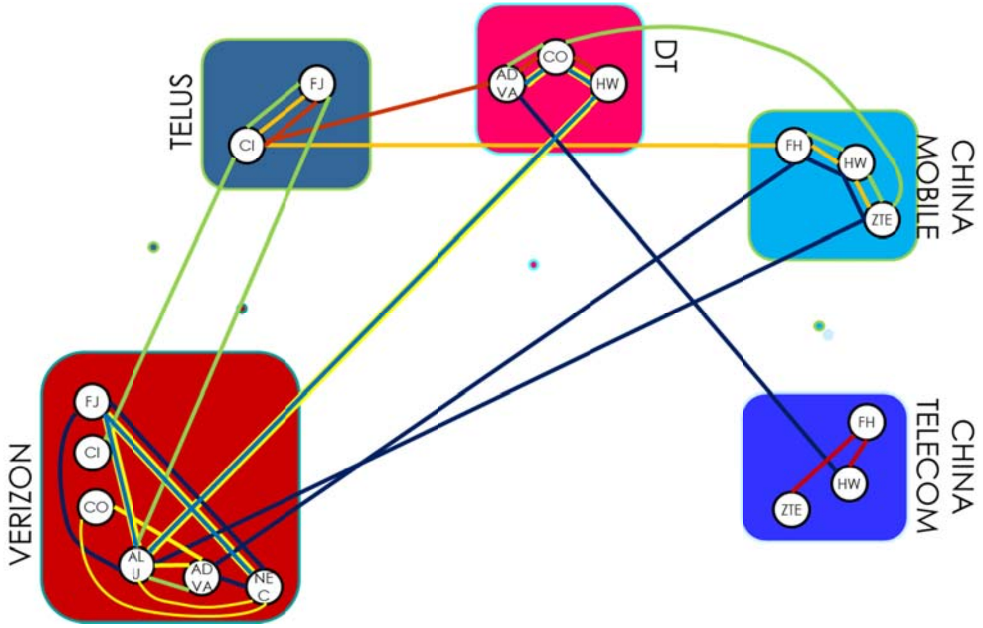

Figure 3: Worldwide Test Topology

As shown in the figure, testing locations were volunteered by five interested carriers, including carriers in Asia, North America and Europe. The host carriers jointly conducted tests with participating vendors.

Participating Carriers:

- Asia: China Mobile, China Telecom
- North America: TELUS, Verizon
- Europe: Deutsche Telekom

Nine vendors, one research institute and two carriers participated in prototyping, providing a variety of types of equipment and software consistent with the SDN architecture.

Participating Vendors:

- ADVA Optical Networking
- Alcatel-Lucent
- Ciena
- Coriant
- FiberHome
- Fujitsu
- Huawei Technologies Co., Ltd.
- NEC Corporation of America
- ZTE

Consulting Carriers and Research Institutions:

- China Academy of Telecommunications Research
- KDDI R&D Laboratories
- Orange

One of the key characteristics of the demonstration was the ability to test applications, controller implementations and optical network elements implemented by different organizations, interoperating through prototype standard or common interfaces.

Features Tested

Features tested included:

- OpenFlow extensions for the Southbound Interface between Controller and Network Element (CDPI) and between Controller and server Controller (CVNI)
- A prototype Northbound Interface between Application and Controller for exchange of Service Requests and responses
- A prototype Northbound Interface for exchange of Topology Requests and responses

OpenFlow Testing

For the demonstration, SBI testing focused on implementations of OpenFlow with optical extensions as defined by the Optical Transport Working Group of ONF[2]. The goals of testing were to prototype these extensions on optical transport equipment and identify areas where specifications may need clarification or correction based on interoperability testing of multiple implementations.

As the optical extensions are still under discussion within ONF, experimental versions of the extensions were tested using OpenFlow 1.3[3] as the base protocol and using experimental extension mechanisms supported in OF 1.3. Extensions tested included:

- Match extensions to allow matching on optical fields
- Port extensions to allow description of optical ports to the controller

In addition to testing of the CDPI directly between the controller and a network device, testing included implementations of the CVNI between parent and child controllers, where the child controller provides a virtualized network representation to the parent controller, making it possible to hide the details of the underlying network. A potential deployment case for CVNI is the ability to support OpenFlow control across legacy domains that do not support native OpenFlow on the network elements.

---

[2] https://www.opennetworking.org/working-groups/optical-transport

[3] https://www.opennetworking.org/sdn-resources/onf-specifications/openflow

OpenFlow Test Cases

CDPI test cases

The Test Cases used for the CDPI (Controller to Switch) tested the following functionality:

1. Connectivity between the Controller and switch including exchange of OpenFlow messages for session establishment and retrieval of capability information from the switch.
2. Ability of the Controller to use OpenFlow messages to install a match table entry creating a cross connection in the switch, verifying in the data plane that the data is flowing. Connectivity cases included:
    - GE port to GE port using ODU0 and 10GE port to 10GE port using ODU2
    - ODU0 port to ODU0 port and ODU2 port to ODU2 port
    - GE/10GE port to ODU0/ODU2 port and ODU0/ODU2 port to GE/10GE port using preconfigured GFP adaptation
3. Ability of the Controller to use OpenFlow with multiple switches to establish a connection across the domain, verifying in the data plane that dataflow exists
4. Ability of the Controller to delete match table entries to tear down the connection.

CVNI test cases

Abstract switch

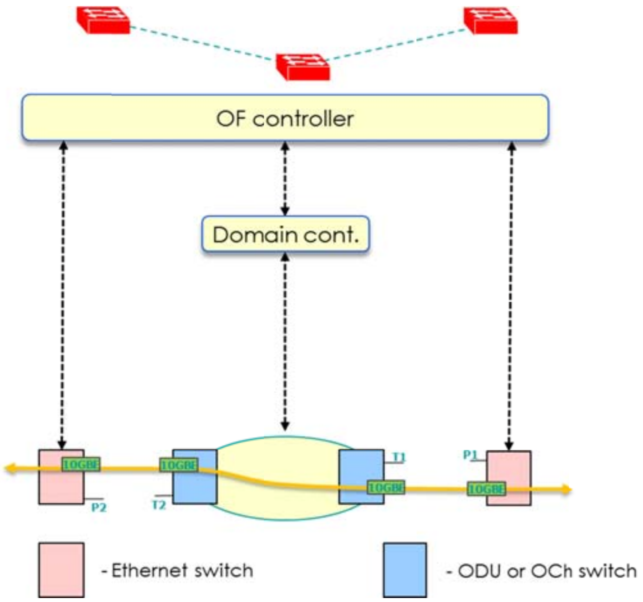

Figure 4: Abstract switch case

In the Abstract switch case, the Parent Controller uses OpenFlow to the Domain Controller as if it is a single switch.

The main change from the CDPI case is that the Domain Controller provides a list of external ports of the domain as if they were ports on a single switch, and responds to flow entries for forwarding between ports by setting up connections between the associated switches. Connectivity cases tested were the same as for the CDPI test cases.

Abstract link

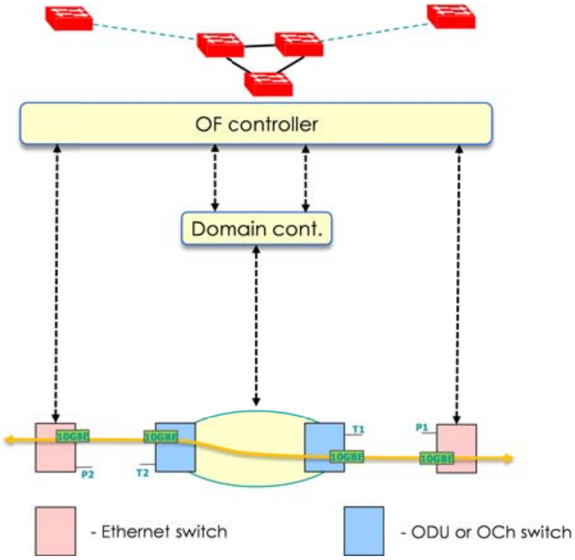

Figure 5: Abstract link case

In the Abstract Link case, the Parent controller uses OpenFlow to talk to a Domain controller, but the Domain controller in this case responds to the Parent controller with an abstract topology consisting of multiple virtual switches connected by links.

Using existing OpenFlow this requires that a separate session be set up between the Parent controller and each virtual switch, and that the Domain controller translate between OpenFlow match entries for each virtual switch and the corresponding actions required in the network.

General Test Result

A total of 8 switch vendors participated successfully in testing of OpenFlow with 4 different controller implementations. The range of optical extensions differed with some vendors implementing only the extensions to the match table to allow OTN connection and other vendors implementing both the match and port extensions.

Both the CDPI and CVNI interfaces were tested. Participation in CVNI testing was of particular interest to China Telecom, which implemented a Parent controller to talk to local Domain controllers. Both CDPI and CVNI interfaces were tested successfully between multiple implementations.

Detailed findings regarding areas of the specifications requiring correction or clarification have been fed back directly to the Open Networking Foundation groups involved in definition of the OpenFlow specification.

Multi-session issue on CVNI

For the CVNI, if the Domain Controller needs to expose a virtual topology with multiple virtual nodes and links to the Parent Controller, the Domain Controller may have to initiate separate OF sessions between the Parent Controller and each virtual switch. This may cause excessive overhead if the Domain Controller is supporting a representation with a large number of virtual switches (for example, it is implemented on the EMS for a large domain of potentially hundreds of switches).

This concern has been raised to ONF and methods of reducing the overhead for the CVNI case are under study.

Service API Testing

The Service API allows the application to request a connectivity service from the network. A common Service API allows a variety of applications to access services provided by the network, especially in an environment with multiple domains with potentially different underlying control methods. In the demonstration, the different domains supported a number of SBIs from the Domain controller, including vendor-specific, standard OpenFlow version 1.3, and OpenFlow with optical extensions as tested in OpenFlow testing described above. The use of a common Service API allowed the same application to be tested across these heterogeneous domains.

The Service API supports four basic test cases:

1. Service creation to create connectivity between endpoints A and Z, which may be in different domains, at a specified layer and bandwidth.
2. Service listing, the ability to request the controller to provide the list of currently active services.
3. Service query, the ability to query the controller about the details of a specific service instance.
4. Service deletion, the ability to terminate connectivity service.

The Service API was designed based on REST and JSON principles to allow for rapid and flexible application design. The API was deliberately simplified for Demo implementation and future work is anticipated to lead to API standards.

Test Configuration

An example of a multi-domain service is shown in Figure 6.

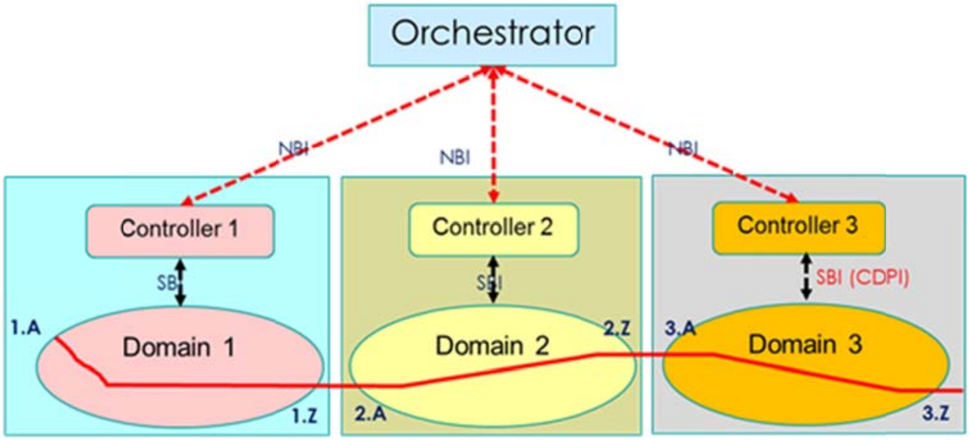

Figure 6: Example of a multi-domain service

For testing of OTN services, the source and the destination nodes were equipped with OTU2 OTN edge interfaces, while for testing of Ethernet services, the source and destination nodes were equipped with 1GbE or 10GbE client service interfaces. GE services were not tested in multi-domain scenario, because the inter-domain links were OTU2, and it was not possible in the simplified API to select a specific timeslot of the inter-domain OTU2 interface.

For intra-carrier testing, the client interfaces were connected to Ethernet or OTN test sets. In one carrier lab, the client interfaces were connected to Ethernet access switches and live non-demand video was streamed across multi-vendor Ethernet and OTN domains. For inter-carrier test cases, the source and the destination client ports were connected to an Ethernet traffic generator or an OTN tester while the E-NNI ports to other carrier labs were looped to simulate an inter-lab physical link.

General Test Results

A total of 7 switch vendors implemented the Service Request API for the demonstration and tested successfully with 4 Network Orchestrator implementations. All 4 base tests were carried out successfully, while a subset of the vendors and orchestrators also implemented and tested an explicit path optional test case. Development and implementation of the Service API as defined for the demonstration was relatively fast due to the choice of REST and JSON as the basis for the API.

The Service Request API allowed orchestrators to create services that spanned multiple heterogeneous domains and multiple carrier lab test sites, using the Demo SCN to carry messaging between different labs. Transit time between labs was not found to be a major factor in supporting the Service Request API.

The Service Request API also allowed multiple orchestrators to access the same set of controllers simultaneously so that each orchestrator could control a subset of the overall demonstration network resources for their own services. This scenario captured many of the features required for the "cloud-burst" application, where different data center operators are able to use their own virtual slices of the service provider network.

Observations of Note

While the Service API successfully supported network orchestration of multiple domains with geographic separation, some aspects needed to be handled through manually intensive means, especially the selection of interdomain links to be used for individual services. Automation of interdomain link discovery and selection would significantly reduce the overhead for using the Service API in a real network with large numbers of switches and interdomain links.

Additionally, the "cloud-burst" scenario was accomplished through manual coordination of the resources assigned to each network orchestrator. A service to be deployed in a real service provider network would need added functionality such as authentication, authorization and scheduling or calendaring of resource requests in order to provide a secure and efficient virtualized network service.

Extensions of the APIs to support such additional functionality will be potential areas of future work in OIF and ONF.

Topology API Testing

The topology API defined the basic topology components in the network and the operations that support the lifecycle of a Transport network. To provide this support, the API enabled the Creation, List, Query, Retrieval and Deletion of topology associated objects. In the demonstration, the List and Read operations were tested.

The API included four topology objects:

1. Vertex: represents a switching node.
2. Edge end: represents an interface on a switching node.
3. Edge: represents a topological link connecting two Edge Ends.
4. Edge end resource: represents the specific resources available at an Edge end.

The response of a List operation returned the list of IDs of the specified topology object (i.e. vertex, edgeEnd, edge, edgeEndResource).

The response of a Read operation returned the ID of specified topology object and its mandatory attributes.

Test cases included:

- Bulk topology transfer
- Specified topology object transfer
    - Specified vertex transfer
    - Specified edge end transfer
    - Specified edge transfer
    - Specified edge end resource transfer
    - Retrieve vertex list
    - Retrieve edge end list
    - Retrieve edge list
- Topology Visualization
- Path Computation

General Results

A total of 5 vendors implemented the Topology API and tested with 2 orchestrators implementing the API. All 10 test cases were carried out successfully. 3 vendors implemented the optional edge end resource object and tested successfully as well.

Observations

The Topology API can support all topology object visualization and path computation across multiple domains successfully.

Benefits

In general these multi-vendor tests conducted jointly with carriers in their labs provide several benefits to participating OIF and ONF members:

- Carriers influence the features and requirements of technology and get equipment that meets their needs
- Interoperability of features across multiple vendors allows carriers to deploy services more rapidly
- Carriers have the opportunity to test vendor interoperability and equipment first hand
- Vendors lower their risk of development because of common functionality, design and component characteristics
- Vendors have a neutral ground to test implementations against others for interoperability and improve their implementations

SDN and virtualization promise to simplify optical transport network control by adding management flexibility and programmatic network element control to enable the rapid services development and provisioning. Improved network efficiency and agility will likewise deliver benefits of lower overall operational expenses and faster time-to-market/revenue resulting in improved ROI for carriers and operators. To this end, participating carriers and vendors leverage the prototype demo to gain practical experience with Transport SDN technology in real-world scenarios in order to assess the status of the technology, develop pertinent use cases, and identify any interoperability and operational challenges that may slow the evolution to commercial deployments. The multi-vendor nature of the testing performed in carrier labs gives carriers the confidence that different transport vendors/systems can work together.

SDN Benefits for Carriers/Operators

SDN is characterized by a physical separation of the forwarding layer and its control plane and the use of a standardized protocol between them, thus enabling the network control to run on a commodity compute engine and in a centralized fashion. This paradigm bypasses the proprietary implementations of traditional network elements and opens the door for interoperability between forwarding layers and control planes provided by different vendors. Further, this migration to an open architecture supports function and feature implementation on common-off-the-shelf (COTS) hardware. As well, it introduces the ability to implement significant improvements in network efficiency, resiliency and new services.

Regular patterns of time of day and day of week usage are seen for specific classes of users and access networks. For example; enterprise users generate most traffic during weekdays and normal business hours while on the other hand consumers typically generate most traffic during nights and weekends. This means that time of day sharing between enterprise and consumer usage patterns will be possible. Rearranging optical network topologies to interconnect networking and computing more economically based upon time of day and day of week to better serve these predictable phases of behaviors could significantly reduce overall networking cost and create new service opportunities. Furthermore, operators may improve the resiliency of their network cores through coordinated multi-layer restoration techniques implemented through a centralized network control function that includes a view of both packet and optical layer topologies and the ability to steer the re-allocation of resources in response to both on-network and off-network failures. This capability may deliver significant improvements in network traffic engineering efficiency and disaster recovery response. For example, the loss of a peering connection to a content provider in one data center could be mitigated through a coordinated reallocation of routing and optical transport resources to geographically diverse peering points.

SDN Benefits for Vendors

As it has been stated already, network operators are currently facing major challenges in order to address not only a huge increase in the overall Internet traffic volume but also a significant change in its nature. The reasons behind this include but are not limited to: the prevalence of the cloud computing paradigm and the associated inter-DCN traffic boom, the rapid increase of mobile user traffic due to the success of smart devices as well as the huge amounts of IP-based video transmissions at ever-improving qualities (from standard HD to 4k video and so on). In order to remain relevant and profitable in the emerging landscape, operators are beginning to impose strict requirements in the form of fast and automated service establishment, teardown and restoration, increased utilization of the underlying network infrastructure, effective network virtualization, as well as offering differential treatment to different types of traffic. Until now, such tasks implied cumbersome control and management making it extremely hard, if not impossible, for vendors to come up with satisfactory solutions to fully address them.

This efficiency gap can be significantly overcome through an adoption of SDN-based paradigms from the vendors regarding their transport products. By cutting down the current costs associated with keeping up with a huge number of supported protocols and the development of elaborate custom APIs for their products, vendors would instead have the opportunity to focus on offering complete end-to-end SDN-based solutions leveraging the omnipresence of SDN and the existence of commonly used APIs. At the same time, vendors can more easily exploit synergies between devices that operate in different network layers (e.g. L2/L3 switches and L1 OTN switches) to offer competitive services that leverage cross-layer optimization. Prominent examples include automated bandwidth on-demand services across domains and layers (e.g. similar to the ones showcased by the OIF/ONF demo), fast service restoration owing to a global knowledge of network resources, or intelligent IP traffic offloading through optical links. The establishment of common southbound and northbound interfaces will allow for more effective and efficient deployment of new products and revenue sources.

Presently there is the commonly held notion that SDN - with its promise of "plain vanilla" network devices - will eventually lead to a significant shrinkage of the market for equipment vendors. This is a misleading and superficial approach to the issue; with more careful examination one would at worst expect just a relocation of the vendor competition arena. The success of vendors will still rely on the reliability of their offered products, perhaps even more so since the gradual shift towards more centralized network control architectures makes the network more vulnerable to single-points-of-failure. Instead of spending effort to support an endlessly increasing number of diverse protocols, resulting in more complex and fault-prone products, they will now need to be able to compete in terms of features such as fast response and scalability (e.g. support of an increased number of flows). Furthermore, new opportunities appear in the horizon, since it is even possible that some vendors may enter the newly formed SDN controller and/or "network application" markets, leveraging their deep expertise of the network operation intricacies. Finally, a prevalence of the SDN concept would allow new and smaller vendors to penetrate into certain market segments more easily, because interoperability across multi-vendor domains would be made significantly easier, thus making the operators' vendor selection choices more flexible.

Conclusion

Participating OIF and ONF member carriers and system vendors joined forces to test prototype transport SDN technology in real-world scenarios. OpenFlow™ extensions developed in the ONF Optical Transport Working Group were prototyped and tested in the demo in both CDPI and CVNI forms. Vendors also tested prototypes of Controller Northbound interfaces for Service Request and Topology functions in development by the OIF. The framework of the demo was cloud-bursting or application-based bandwidth-on-demand between data center sites. Testing based on this real-world use case illustrates potential deployment of Transport SDN technology, common interfaces required, needs for interoperability and any operational challenges.

The experiences from the demo will be shared across the industry to help develop Implementation Agreements and specifications. The ultimate goal is to accelerate the deployment of practical, programmable transport networks that enable a new era of dynamic services.

Appendix A: List of Contributors

Dave Brown (Editor) – Alcatel-Lucent
Hui Ding – China Academy of Telecommunications Research
Hans-Martin Foisel – Deutsche Telekom
Jia He – Huawei Technologies Co., Ltd.
Konstantinos Kanonakis – NEC Laboratories America, Inc.
Rob Keates – TELUS
Yunbo Li – China Mobile
John McDonough – NEC Corporation of America
Lyndon Ong – Ciena
Karthik Sethuraman – NEC Corporation of America
Vishnu Shukla - Verizon

Appendix B: About the OIF

Launched in 1998, the OIF is the first industry group to unite representatives from data and optical networking disciplines, including many of the world's leading carriers, component manufacturers and system vendors. The OIF promotes the development and deployment of interoperable networking solutions and services through the creation of Implementation Agreements (IAs) for optical, interconnect, network processing, component and networking systems technologies. The OIF actively supports and extends the work of standards bodies and industry forums with the goal of promoting worldwide compatibility of optical internetworking products. Information on the OIF can be found at http://www.oiforum.com.

Appendix C: About the ONF

Launched in 2011 by Deutsche Telekom, Facebook, Google, Microsoft, Verizon, and Yahoo!, the Open Networking Foundation (ONF) is a growing nonprofit organization with more than 140 members whose mission is to accelerate the adoption of open SDN. ONF promotes open SDN and OpenFlow technologies and standards while fostering a vibrant market of products, services, applications, customers, and users. For further details visit the ONF website at: http://www.opennetworking.org.

Appendix D: Glossary

API	Application Programming Interface
CDPI	Control to Data Plane Interface
COTS	Common-off-the-shelf
CVNI	Control Virtual Network Interface
DCN	Data Communication Networks
E-NNI	External Network-Network Interface
GFP	Generic Framing Procedure
IA	Implementation Agreement
JSON	JavaScript Object Notation
NBI	Northbound Interface
ODU	Optical channel Data Unit
OF	OpenFlow
OIF	Optical Internetworking Forum
ONF	Open Networking Foundation
OpEx	Operational Expenditure
OTN	Optical Transport Networking
OTU	Optical channel Transport Unit
QoE	Quality of Experience
QoS	Quality of Service
REST	Representational state transfer
ROI	Return on Investment
SBI	Southbound Interface
SDN	Software-Defined Networking
TCO	Total Cost of Ownership