Demystifying SDN for Optical Transport Networks: Real-World Deployments and Insights

Abhinava Sadasivarao, Dharmendra Naik, Chris Liou, Sharfuddin Syed and Anurag Sharma
Infinera Corporation, Sunnyvale, CA, USA
{*asadasivarao, dnaik, cliou, ssyed, ansharma*}@*infinera.com*

***Abstract*—We present our observations deploying SDN-enabled solutions in operational Tier-1 carrier networks. This includes technical and operational issues that are of prime importance to service providers. Specifically, we focus on two archetypal customer use cases that we have observed from the point of view of optical transport applications. In addition, other key aspects such as the role of embedded control plane in SDN are discussed. Development models used by service providers and how this plays a crucial/symbiotic role between the vendor (enhancing platform data & control plane features) and the service provider (rolling out SDN-enabled services) are considered. We conclude with a synopsis on the state-of-affairs, and the role of standardization in the SDN ecosystem.**

1. Introduction

The evolution (and adoption) of Software Defined Networking (SDN) concepts, techniques and applications has been slow but steady. Service providers across enterprise, data center and WAN segments are actively engaging with vendors to explore possibilities to improve network operational efficiency - all the *-ilities* such as programmability, manageability, mobility & evolvability. Large Internet Content Providers (ICP) have demonstrated the necessity for programmable abstractions given the petabyte scale, data intensive applications for which a robust network infrastructure becomes fundamental [1], [2], [3].

In the recent past, the primary focus of SDN within the industry (as well as academia) has centered around Layer 2, Layer 3 and above. The underlying Layer 1 and Layer 0 networks have often been considered "fixed" in the sense that, their primary responsibility is to act as a point-to-point underlay, providing the necessary capacity between Ethernet/IP domains. Advancement in optical technologies, bit rates and deployment of Optical Transport Network (OTN) [4] have enabled transport networks to provide flexible grooming, multiplexing and switching functions in addition to data transport and survivability. Transport network elements are being supplemented with more intelligent *"flow aware"* and *"packet aware"* features for improved visibility. The growth in traffic volumes, dynamic traffic profiles and types of applications requires service providers to give careful consideration to the optical transport backbone.

Even in data center networks, where flows are expected to be highly dynamic and optical switching has not been traditionally used, there are proposals that explore/advocate architectures which involve optical switching to scale efficiently without losing flexibility; enabled by the advances in photonics and electronics in recent years [5], [6].

We envisioned Open Transport Switch virtual (OTSv) as a programmable architecture for transport networks [7]. Several pilot experiments were conducted with carriers as well as Research & Educational partners to validate key OTSv use cases. These include multi-layer visibility and WAN optimization [8], unified packet-optical control plane [9], multi-vendor optical control plane integration [10] and data center optical interconnects [11].

Our goal in this paper is to provide perspectives deploying SDN applications in *operational* optical transport networks, primarily in the WAN context. We present our experiences deploying OTSv in Tier-1 carrier networks, along with two key use cases for which, SDN techniques would be a good fit: Layer 1 Virtual Private Networks and Port/flow services. We also present our views on other critical aspects such as the role of embedded control plane in an SDN environment and integration/development models used by carriers (DevOps v/s NetOps). We conclude with a brief summary on the ongoing standardization efforts, lessons learned and future outlook.

2. Overview of OTSv

This section provides a high-level overview of OTSv. Although the architecture of OTSv itself is not our focus, a brief overview would set the stage for subsequent discussions.

The architecture originally proposed in [7] provided a nodal abstraction of a typical Layer 1 optical transport Network Element (NE). The primary programmable interface was OpenFlow v1.0 [12] and relied on vendor extensions for optical capability discovery. Optical constructs such as time-slots and cross-connects (XCON) were introduced as OpenFlow primitives. These extensions catered to providing both models of programming viz. *explicit* (where the controller layer is aware of every Node between service end-points and programs forwarding at every hop) and *implicit* (where the controller layer is only provided visibility of edges and embedded nodal Control Plane computes the path between the end-points). Although the choice of OpenFlow was to ease integration with commonly available (popular) controller platforms, the abstractions were protocol agnostic. We have, since then, enhanced OTSv in two aspects: (i) Provide domain & network abstractions (in addition to nodal abstraction) (ii) Support for other southbound NE interfaces such as NETCONF [13].

The OpenFlow implementation has been enhanced to utilize ONF Optical Transport Working Group (OTWG) extensions [14] (which is a work-in-progress to embed optical primitives natively into OpenFlow).

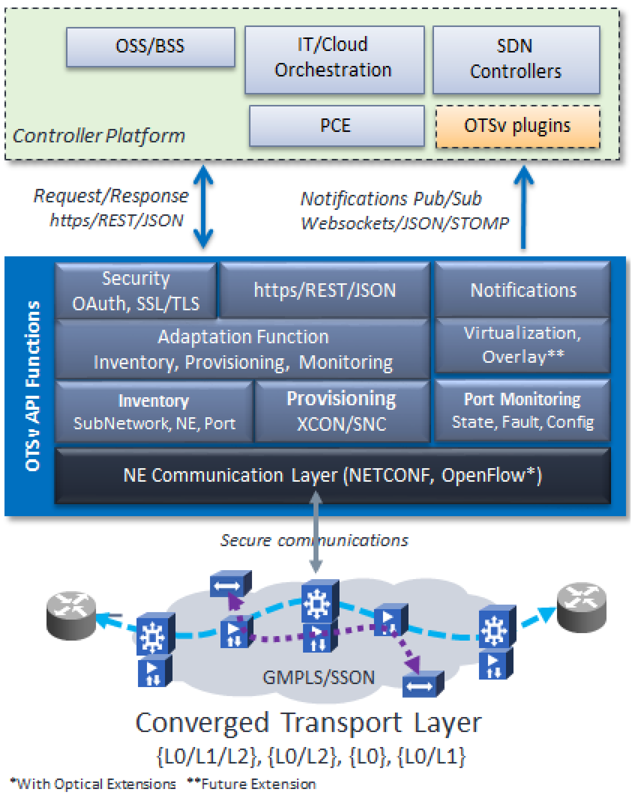

Figure 1. Overview of OTSv Internals

Fig. 1 gives an overview of the OTSv internals. The system doesn't inherently provide any path computation capabilities; it provides full network visibility allowing the controllers to implement desirable service objectives. Customers implement their own Path Computation Element (PCE) functions which calculate suitable path(s) for service provisioning. Typical routing metrics include SRLG diversity, shortest path based on TE-Link administrative costs, source-to-destination hops and latency.

One of the most important ingredients is the information model. We use the IETF multi-layer data model proposed in [15] for the Traffic Engineering Database (TED). The model includes vendor-agnostic topological representations for TE-Links, neighbor adjacency and bandwidth accounting. Optical entities such as OTN Optical Channel Data Units (ODU) time-slots, Layer 0 frequency slots/spectrum (carrier width and center frequency) are modeled as first-class primitives. The model also provides the ability to express nodal switching matrix/constraints.

We use ONF OTWG information model to represent managed entities such as node, sub-network, domain (forwarding domain), flows (forwarding constructs).

Other OTSv features include XCON time-slot creation & retrieval, creation of GMPLS [16] Signaled Subnetwork Connections (SNC), fault monitoring at various digital layers (Client layers - SONET/SDH, Ethernet and Server layers ODU, OTU) and optical layers (Optical λ and Optical Carrier Groups).

3. Optical Transport SDN Use Cases

3.1. Port and Flow Services

The most common scenario involves provisioning Layer 1 transport circuits. The client signal (payload) to be transported (Ethernet, SONET/SDH, FibreChannel etc.) is mapped to a right-sized ODUk (ODU0 = 1.25G, ODU1 = 2.5G, ODU2/ODU2e/ODU1e = 10G, ODU3 = 40G, ODU4 = 100G), switched within the OTN domain and eventually, the client signal is dropped at the destination. Optical to Electrical (O-E-O) conversions occur along intermediate hops in the optical (and optionally digital) domain for Reamplify - Reshape - Retime (3R) regeneration depending on the reach.

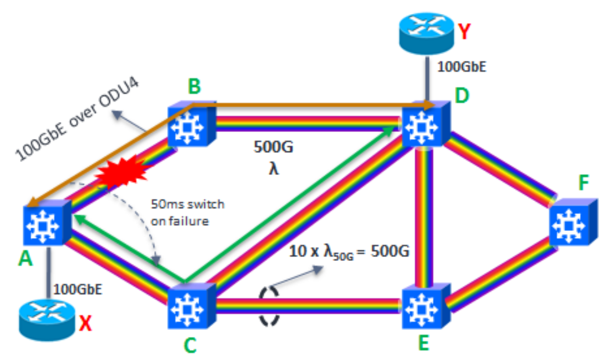

Figure 2. 100GbE over ODU4 Transport Example

Fig. 2 shows such a port service example where 100G Ethernet transport service is created between routers X and Y. An ODU4 service is setup between Node A and Node D. The client 100GbE is "added" at A and "dropped" at D. The ODU4 is switched in the OTN domain at intermediate node B. The setup of this Layer 1 service could be:

- Through XCONs where the SDN controller has visibility of the Layer 1 topology and computes the end-to-end path. XCONs are created between input and output time-slots at each hop at A, B & D
- Through signaled GMPLS SNC (similar to an MPLS LSP). Here, the controller's visibility into the optical domain is only at the edges. An SNC request is made to the head-end (Node A) or tail-end (Node D) to create a bidirectional 100GbE circuit. The node first performs OSPF-TE route query to compute a 100GbE path (based on shortest path or latency) and subsequently issues an RSVP-TE request to signal the SNC

If customer SLA indicates recovery from failures, the service is "marked" for protection (or restoration) where upon detecting failure, sub-50ms protection switching is performed to restore the service. In Fig. 2, if the active Layer 1 path A $\leftrightarrow$ B $\leftrightarrow$ D fails, protect path A $\leftrightarrow$ C $\leftrightarrow$ D (pre-established or dynamically computed) can be activated to restore the service.

On the other hand, flow (sub-port) services are mainly multiplexing services which utilize the OTN multiplexing hierarchy to carry higher-order ODU (HO-ODUk) which could contain several lower-order ODU (LO-ODUj). The OTN containers can also be composed of dynamically resizeable ODUflex with ODU0 (1.25Gbps) granularity. Ethernet services (port and VLAN) of different rates get mapped to these LO-ODUj. The HO-ODUk instantiation can be done separately from the LO-ODUj at different points in time.

3.2. Layer 1 Virtual Private Networks

Alternately called *Layer 1 Virtualization*, is the ability to partition an existing optical infrastructure into multiple, logical virtual transport networks. These partitions could be on a per-customer basis. Once partitioned, these slices are sandboxed with the customers being able to fully control their respective overlays.

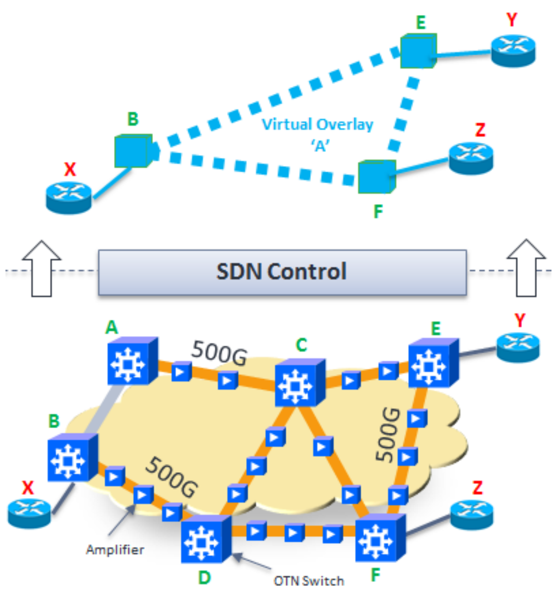

Figure 3. Transport Network Virtualization

Fig. 3 shows an example where a virtual transport overlay ('A') has been sliced from the underlying physical network. The customer partition only has port visibility at the edges (B, E & F) where the customer edge routers (X, Y & Z) are collocated. Several overlays could exist per customer making the physical infrastructure multi-tenant.

The Transport-aware SDN controller presents this *"virtualized view"* of the network to the respective tenants. The realization of resource partitioning on the device *could also* rely on an embedded control plane (if available) to create such partitions (For e.g., use of OSPF affinities for TE partitioning [17]). There exist multiple schemes that can be used for network partitioning; Loosely constrained (where slices are created by specifying coarse thresholds, FCFS resource assignment) or strictly constrained (hard resource partitioning) are some examples but these are outside the scope of this paper.

4. Operational Deployment

One of OTSv deployments happened in March 2015 with PACNET [18] (now part of Telstra [19]). Trial runs on isolated network segments were conducted before going live on PACNET's operational network. PACNET is one of the larger telecom providers in the APAC/Oceania region. PACNET/Telstra operates terrestrial and submarine networks, catering to both consumer and enterprise customers.

PACNET Enabled Network (PEN) is a PACNET networking platform to develop (EPL/EVPL Ethernet and IP/MPLS VPN) services using a variety of software and hardware components. This includes traditional routers (including LER & LSR), optical transport switches and *NFV* devices. The optical transport devices are comprised of Infinera DTN-X platform [20]. Some noteworthy aspects are:

- Operational Model: PACNET acts as the provider and their customers (Tier-2's or other enterprises operating metro & data center networks) act as consumers. These consumers *bring their devices* to PACNET's PoP for port attachment.
- Some of the SLA/QoS metrics available to customers include low latency, dedicated v/s best effort, pay-per-bit and pay-as-you-go. Billing is based on time-of-day and/or bandwidth consumption.
- PEN platform can alternatively be deemed as *"controller Layer"*. A customer intent of *"Setup an Ethernet Private Line (EPL) service from Sydney to Tokyo with CIR of 10Gbps"* is translated into multiple steps, percolating through the controller eventually as a set of forwarding/flow instructions to the networking devices where the service is realized.
- PEN utilizes OTSv for control of Infinera DTN-X [20] transport domain within the overall PACNET network. Other L2/L3 devices and NFV appliances are also integrated into PEN (via OpenFlow, NETCONF and other vendor interfaces) for end-to-end provisioning and monitoring.

PEN PoPs are across the globe in 13 cities and 8 countries [21], which are mostly data centers. Additionally, there are also various "Cloud Endpoint" sites where PEN allows customers to interconnect with public clouds such as Amazon AWS. Fig. 4 shows a segment of the global PEN sites in the APAC/Oceania region (pacific rim).

Those sites interconnected by the terrestrial and submarine links **highlighted** in Fig. 4 have DTN-X optical transport network elements which are controlled by PEN via OTSv. Ethernet 1/10/100Gbps over Layer 1 OTN circuits

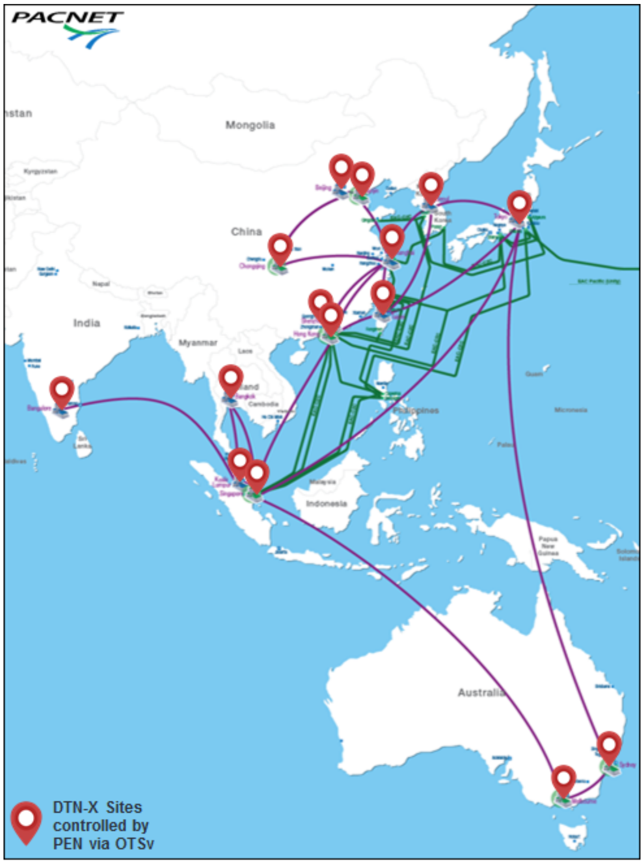

Figure 4. Overview of PACNET/Telstra OTSv Deployment

are provisioned (either dynamically or pre-provisioned) on these fiber spans allowing turn-up of hundreds of Gigabits of underlay for customer Ethernet interconnection. PACNET operations' feedback indicated 10GbE over ODU2e as the most commonly provisioned L1 optical service. As of this writing, PACNET has commercialized PEN for service delivery.

5. Discussion

5.1. Key Observations

The PACNET example is one of the optical SDN deployments where transport network elements were integrated into the provider's SDN infrastructure. The integration of optical transport elements into the "overall network" for the purposes of automation and monitoring/diagnostics & debugging is the *primary* driver for transport layer programmatic capabilities. We observe a growing pattern emerging with other provider deployments with the following highlights:

- Most common optical SDN deployments at the moment are for traditional port-based services. Flow (sub-port) services are expected to fast-follow port services given the burstiness/granularity of Ethernet services. ODU0 granular OTN switching at Layer 1 would be *fundamental* to avoid stranded bandwidth fragmentation. Providers wish to make *"at what layer to switch/groom a service?"* decisions dynamically.
- Use of GMPLS Control Plane: Majority of field deployments as well as trials seem to indicate the use of Signaled SNC over explicit XCON creation. Although OTSv provides all the necessary information on the underlying topology, bandwidth/timeslot availability that would allow providers to perform path computation in their controller layers, traction has been towards the use of embedded control plane. We provide some (possible) insights on why this could be the case in Section 5.2.
- L0 SDN control of ROADM configurations & impairment aware wavelength routing is going to be important. However, this will likely be of interest (in the next ~2-3 years) to very large Tier-1 providers and ICPs *only* who operate L0 networks. Given the growing capacities, the emphasis is on *"Pack as many Optical λ on the fiber as possible"* (subject to physical fiber characteristics and other analog parameters like modulation such as QPSK, 8QAM, 16QAM etc.)
- NE interfaces to communicate with the devices need to be *"friendly"* but not necessarily be standards compliant (like OpenFlow). However, providers give careful consideration to the underlying data & information models for device, service, topology and bandwidth representation. Given that the ecosystem is still in its nascent stages, the focus on APIs is centered around *"provide APIs for the most common use cases"* rather than exposing each and every artifact that can be controlled on the device.
- *Integration has challenges*; An initial learning curve, hand-holding customer personnel in helping them understand APIs, abstractions and information models. Operational success is largely dependent upon equipment vendor $\leftrightarrow$ carrier interaction. Rather than the carrier having to wait for longer periods for vendor feature delivery (subsequently followed by extensive certification to perform OSS integration), the vendor enhancements have to be incremental, with iterative faster turnaround times.

5.2. Is There a Role for Control Plane in SDN?

One of the central tenets of SDN is data and control plane separation. The idea is for equipment vendors to provide programmable interfaces allowing routing, QoS decisions to be made outside the device. One of the primary reasons is to enable faster innovation (not locked to vendor development cycle). While these remain true, some of these drivers slightly differ when it comes to optical transport networks.

Although standards exist for data plane interoperability, Layer 1 data plane interop has challenges (as opposed to

Ethernet/IP). Every vendor has proprietary frame formats (in the interior network domain/NNI). In addition, the granularity of grooming functions differ, making it non-optimal to define generic, vendor-agnostic models for timeslots, connections, multiplexing identifiers, topology discovery and other traffic carrying entities. This problem is compounded further for Layer 0 (optical domain) due to proprietary modulation/encoding and Forward Error Correction (FEC) schemes. Protection/restoration add additional complexities; given the stringent requirements on protection (typically sub-50ms), a traffic outage like a fiber cut can impact tens/hundreds of gigabits. Hence, it is best for performance reasons if protection/restoration functions are left to the embedded control plane.

There also exists an interesting hybrid approach where the end-to-end Layer 1 path is computed in the controller, but relies on the RSVP-TE component of the control plane to signal the path. The SNC creation request has the full explicit route object (ERO) inclusion list. This provides the best of both worlds as the burden of LSP setup/teardown/crankback is handled by the device's RSVP-TE allowing the controller to focus on routing aspects if path computation needs to be centralized.

For example, in Fig. 2, the SDN controller could compute a Layer 1 OTN path from C to F, followed by a GMPLS request to either C or F to establish an SNC (with the full ERO: $\{C \leftrightarrow E \leftrightarrow F\}$ or $\{C \leftrightarrow E \leftrightarrow D \leftrightarrow F\}$)). RSVP-TE ensures the SNC creation, ensuring transactional guarantees, returning `SUCCESS` or `FAILURE` indication.

OTSv allows both manual XCON creation as well as GMPLS control plane SNC. Due to the reasons described, we observe customers preferring signaled SNC over XCON provisioning. We expect embedded GMPLS to be in use, at least in the near future.

5.3. Multi-Domain, Multi-Layer and Multi-Vendor Orchestration

Another recurring point of discussion is centered around multi-layer and multi-domain orchestration: *The ability for a service provider to have a homogeneous, uniform, consistent view of the entire network*. Large Tier-1 networks are inherently multi-vendor; The carriers (for various business reasons) have multiple vendors from whom they procure networking equipment. Historically (as things stand today), the carrier operational personnel partition the control and management planes into multiple domains for ease of OAM&P. But given the differences in the NE capabilities between vendors, seamless network visibility comes at a cost. For example, North American Tier-1 carriers (and their equipment vendors) invest significant resources in OSMINE certification [22] where a trusted 3rd party (Telcordia) performs multi-vendor management and control plane integration for the carrier. Certification processes like OSMINE are adopted primarily by large carriers (also restricted by geographic regions) and the cost/time investment make them nonviable for Tier-2 and smaller carriers

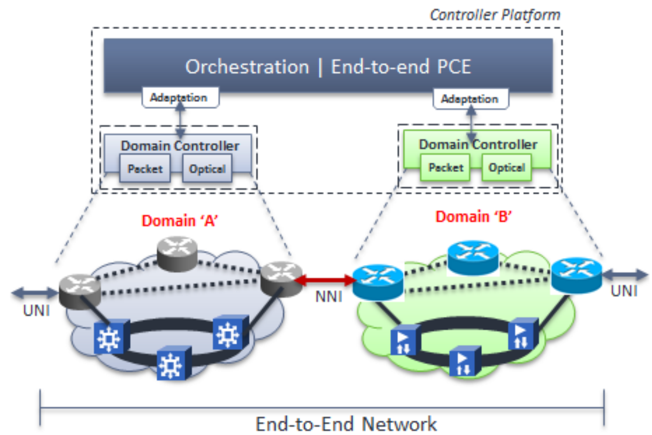

Figure 5. Multi-Layer and Multi-Domain Coordination

Layer 2 and Layer 3 networks have traditionally been interoperable both in data (E.g.: IP/MPLS) and control planes (E.g.: BGP, OSPF). As far as Layer 1/Layer 0 optical networks are concerned, data plane interoperability is in theory possible; ITU specifications [4], [23] allows for multi-vendor digital and optical interface interworking. L1/L0 optical control plane interop, however, has a lot of challenges. GMPLS control plane implementations across optical vendors seldom interoperate; Alternate standardization efforts haven't seen wide adoption [24], [25].

An SDN-enabled network needs to address the multi-layer and multi-domain challenges. End-to-end path computation and provisioning across disjoint administrative domains within a carrier's network requires a uniform view of all the resources in the network. Our experience so far, suggests that the current approach towards multi-domain/multi-layer integration is to introduce hierarchy; The carrier partitions network segments into multiple domains and have *domain controllers* to manage these islands.

Fig. 5 depicts an example. Here, the overall network of a carrier is segmented into two domains ('A' & 'B'). Every domain has a mix of L2/L3 and optical transport elements. Individual *domain controllers* are responsible for controlling the devices/resources within the segments. A higher-level Orchestrator assimilates network state from individual domains to construct end-to-end controls. Adaptation functions ensure that the network information is normalized across different vendors equipment.

5.4. Dynamic Restoration and Traffic Optimization

We briefly discussed optical protection/restoration in Section 5.2. Even if protection functions are left to the NE, there are tangible benefits in providing dynamic restoration. Restoration would involve the SDN controller computing alternate path(s) dynamically, triggered by failures in the network. Multiple options exist here. For example, one approach involves an adaptive scheme where the SDN controller is aware of the probabilities of failures (based on historical alarms/faults or susceptibility to fiber cuts at specific geographical locations). For such network segments, the operator can dedicate sufficient protection during deployment. On multiple failures (beyond the amount of bandwidth reserved for protection), the controller then dynamically recomputes one (or more) paths to restore, when it receives failure indication. Alternatively, the controller can also *continuously* compute alternate paths as and when failures occur: The controller re-optimizes these pre-computed paths in accordance with the updated network state.

It's important to note that restoration is dependent upon the communication latency between the NE and the controller (in addition to the complexity of the path [re-]computation algorithms): The trade-off being the amount of service downtime v/s controller availability (replicated, highly available controller instances). Integration of planning capabilities with the controller adds additional *awareness* to the SDN controller. For example, [26] explores integration of offline and online routing planning into the SDN framework.

Unified multi-layer visibility allows [re]routing decisions to be optimized; For example, knowledge of multiple IP/MPLS links sharing one more optical SRLGs would allow servicing requests which have specific SLA (dedicated, unprotected, best-effort etc.).

5.5. Service Provider Operational Models

The true usefulness of SDN for service provider adoption is critically dependent upon several non-technical aspects. In particular, a lot depends upon the organizational structure of the provider. The traditional model (NetOps) in large service provider organizations includes several groups: one to manage packet networks, another group maintaining transport layers, group of technicians for debugging and so on. Very likely, the business decisions on vendor selection are influenced by these groups (independent of other groups) given the skillset of the support organization (being adept at configuring & debugging a certain vendor's equipment over others).

More recently, *DevOps* models are being experimented upon where providers employ operational personnel with software development skills. But this needs organizational restructuring (significant in cases) to bear fruits. In our PACNET experience, the provider was willing to invest resources as they see long term benefits. On the other hand, other customers are still tentative. There is also a tussle between DevOps and NetOps groups; When it comes to revenues, majority of provider revenues (at the moment) are still due to traditional (L2VPN, L3VPN, OTN Circuits etc.) services (managed by NetOps) and a very small % is due to SDN-enabled services (managed by DevOps). Not to mention, the differences in mindset (risk-averse v/s risk-takers, standards-oriented v/s quick-and-dirty).

More providers migrating to DevOps model will be beneficial for SDN adoption. The vendor $\leftrightarrow$ provider interaction cycles need to be shorter and iterative where provider integration of features follows an agile model: prototype $\rightarrow$ simulate $\rightarrow$ test $\rightarrow$ deploy $\rightarrow$ evolve.

5.6. SDN and Standardization

Standardization efforts in the transport domain (and SDN in general) has been a work-in-progress in different SDOs (ONF, IETF, MEF). These efforts need to continue with concerted involvement from both the vendor and provider communities. It is encouraging to see efforts such as OpenConfig [27] driven largely by providers (traditional carriers & ICPs) where device and vendor-agnostic information models are being proposed. Community driven efforts such as OpenDaylight are providing valuable tools, not to mention, alternative forums outside standards for hands-on collaboration.

Standardization efforts in the transport SDN space, has been (somewhat) sluggish. This is understandable given fundamental differences between transport devices across vendors (data plane incompatibilities, topology modeling etc.). Given that provider deployments have been incremental, the emphasis on standardized interfaces has been moderate.

To the best of our knowledge, OTSv is one of the first implementations to adopt the ONF OTWG information model (albeit being a work-in-progress). Our approach was to keep the architecture as generic as possible from Day #1 to facilitate easier integration into the provider's controller ecosystem. Feedback on the architecture and our general approach to Transport SDN has been very encouraging but as it stands, standards compliance is not make-or-break when it comes to provider adoption.

5.6.1. Optical Transport Service and Policy Definitions. One aspect which would significantly benefit the optical transport community is standardization of L1/L0 services. One possibility is to follow an approach similar to MEF who defines Ethernet services which are in wide adoption today [28]. Similarly, a common, vendor-independent service definitions can be established for typical L1/L0 digital and optical services. This would include common information models and management interface specifications to cover the full spectrum of FCAPS and OAM&P. This would ensure that multi-vendor integration becomes easier and helps with interop testing. Adaptation functions/vendor plugins (as seen in Section 5.3) can (potentially) be reused by the operator across different vendor equipment (which provide similar capabilities).

It goes without saying that care *must* be taken to ensure that any standardization activity does not result in (yet another) plethora of recommendations/guidelines which aren't adopted eventually. As we saw above, carriers and service providers driving/initiating efforts within SDOs will be one way to make sure the use cases are well defined (justifying standardization). Recent initiatives such as [29] are steps in the right direction with respect to optical interoperability and unified management.

5.7. Other Points to Consider

Support for Brownfield Networks: With increasing support for SDN interfaces on modern day networking devices, carriers need a migration path for legacy/existing brownfield networks. This includes network elements as well as network management (NMS/EMS) installations.

Devices which offer embedded interfaces can be brought under the purview of the SDN controller by adapting management protocols (like SNMP or TL-1) for control. Alternatively, the controller can be augmented with plug-ins to interop with NMS/EMS (over interfaces like MTOSI [30]).

Maintenance and Debugging: refers to "*In a multi-layer, multi-domain network, how easy is it to fingerpoint where the service intent is broken?*". The concerns are understandable given that the amount of software systems deployed in an SDN environment is significant (if not greater than traditional network operations/OSS). In a multi-vendor setting, debugging can be fairly complex. Current deployments of (migration towards) SDN have been gradual. Once carriers contemplate a large-scale migration, maintenance issues will certainly magnify. Standardized equipment/inventory, service and topology models homogenize network resources across vendors that are visible to the controller. [31], [32], [33] propose OpenFlow-centric *single-stepping, backtrace & breakpoint* style debugging. Further investigation is needed in developing generic, non-OpenFlow debugging mechanisms.

6. Conclusion

Although the discussion was centered around optical transport, most of the aspects mentioned equally apply to packet transport technologies (MPLS, MPLS-TP). We highlighted our experiences deploying software-defined L1/L0 network controls and a few critical issues that large Tier-1 carriers face - both technical and operational. Some of these challenges can be solved through collaborative standardization efforts (with equal participation from provider and vendor communities). Other issues (such as multi-layer, multi-domain control) have seemingly feasible solutions today but could require substantial re-engineering at larger scales.

SDN for optical transport is going to be an integral part of provider infrastructure. Transport SDN adoptions are still in nascent stages but we believe this will steadily grow over the next 2-3 years. Visibility into, and programmability of L1/L0 layers will be imperative from both technical and operational perspectives. DevOps models and shorter vendor $\leftrightarrow$ provider cycles will expedite SDN evolution. We expect traditional embedded control plane to be in use while carriers devise strategies for a seamless migration to SDN-based network operations, management and control.

References

[1] S. Jain *et al.*, "B4: Experience with a Globally-deployed Software Defined WAN," in *Proceedings of ACM SIGCOMM*, Aug 2013.

[2] A. Roy *et al.*, "Inside the Social Network's (Datacenter) Network," in *Proceedings of ACM SIGCOMM*, Aug 2015.

[3] C.-Y. Hong *et al.*, "Achieving High Utilization with Software-driven WAN," in *Proceedings of ACM SIGCOMM*, Aug 2013.

[4] ITU-T, *G.709: Interfaces for Optical Transport Network*, Std., Feb 2012.

[5] G. Porter *et al.*, "Integrating Microsecond Circuit Switching into the Data Center," in *Proceedings of ACM SIGCOMM*, Aug 2013.

[6] H. Liu *et al.*, "Circuit Switching Under the Radar with REACToR," in *Proceedings of USENIX NSDI*, Apr 2014.

[7] A. Sadasivarao *et al.*, "Open Transport Switch: A Software Defined Networking Architecture for Transport Networks," in *Proceedings of ACM SIGCOMM HotSDN*, Aug 2013.

[8] H. Rodrigues *et al.*, "Traffic Optimization in Multi-layered WANs Using SDN," in *Proceedings of the IEEE Symposium on High-Performance Interconnects (HOTI)*, Aug 2014.

[9] A. Sadasivarao *et al.*, "Enabling Multi-Layer Provisioning and Optimization for Core Transport Networks with Unified Packet-Optical Control Plane," USENIX NSDI '14, Posters & Demos, Apr 2014.

[10] V. Lopez *et al.*, "Demonstration of SDN Orchestration in Optical Multi-vendor Scenarios," in *Optical Fiber Communication Conference (OFC)*, Mar 2015.

[11] A. Sadasivarao *et al.*, "Bursting Data between Data Centers: Case for Transport SDN," in *Proceedings of the IEEE Symposium on High-Performance Interconnects (HOTI)*, Aug 2013.

[12] ONF, *OpenFlow Switch Specifications v1.0*, Std., Dec 2009.

[13] IETF, *RFC6241: Network Configuration Protocol (NETCONF)*, Std., June 2011.

[14] ONF, *Optical Transport Protocol Extensions v1.0*, Std., Mar 2015.

[15] IETF, *draft-liu-yang-ted-00: A YANG Module for Traffic Engineering Database (TED)*, Std., October 2014.

[16] IETF, *RFC3495: Generalized Multi-Protocol Label Switching (GMPLS) Architecture*, Std., January 2009.

[17] ——, *RFC5392: OSPF Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering*, Std., January 2009.

[18] Ovum, "PACNET and Infinera Deploy First Commercial Optical Transport SDN http://www.ovum.com/pacnet-and-infinera-deploy-first-commercial-optical-transport-sdn/," Mar 2015.

[19] Telstra, "www.telstra.com.au."

[20] Infinera, "DTN-X www.infinera.com/products/dtn-x-family/."

[21] Telstra/PACNET, "Global PEN Sites http://www.telstraglobal.com/products/global-connectivity/pen."

[22] Telcordia, "Operations Systems Modifications for the Integration of Network Elements."

[23] ITU-T, *G.694.1: Spectral Grids for WDM Applications: DWDM Frequency Grid*, Std., Feb 2012.

[24] OIF, *External Network-Network Interface (E-NNI) Signaling Specification*, Std.

[25] ——, *User Network Interface (UNI) Signaling Specification*, Std.

[26] Z. Cao *et al.*, "Traffic Steering in Software Defined Networks: Planning and Online Routing," in *Proceedings of ACM SIGCOMM DCC*, Aug 2014.

[27] OpenConfig, "www.openconfig.net."

[28] MEF, *(6.2) EVC Ethernet Service Definitions*, Std.

[29] OpenROADM, "Multi-Source Agreement www.openroadm.org," 2016.

[30] TM Forum, *Multi-Technology Operations Systems Interface (MTOSI)*, Std.

[31] N. Handigol *et al.*, "Where is the Debugger for My Software-defined Network?" in *Proceedings of ACM SIGCOMM HotSDN*, Aug 2012.

[32] R. Durairajan *et al.*, "Controller-agnostic SDN Debugging," in *Proceedings of ACM CoNEXT*, Dec 2014.

[33] A. Wundsam *et al.*, "OFRewind: Enabling Record and Replay Troubleshooting for Networks," in *Proceedings of USENIX ATC*, Jun 2011.