# Optimizing Multi-Layer Networks with Transport SDN

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

**Chris Liou**
Vice President, Network Strategy
Infinera Corp.
169 Java Dr., Sunnyvale, CA 94089
(408)-572-5366
cliou@infinera.com

# Overview

The growth and evolution in traffic types and patterns over IP backbones is causing cable operators and service providers to reexamine the best ways to not only architect their IP and optical transport networks, but also the various ways to manage, orchestrate, and optimize multi-layer networks. The conventional approach is to place all networking functions solely within the IP router layer and to leverage the router vendor's embedded control plane protocols, and employing an optical network to essentially deliver static, dumb "pipes" to interconnect routers. With the recent emergence of SDN technology and its ability to decouple the control plane from the forwarding plane, resulting in direct, explicit control and programmability of packet flows, operators are now beginning to explore whether SDN has a role in improving the operations and economics of their multi-layer backbones and optimizing the performance of their packet networks. Furthermore, as Transport SDN beings to emerge for converged optical systems with integrated OTN/Packet switching, new options for deploying and managing multi-layer, multi-vendor networks emerge.

This paper defines and explores the capabilities that Transport SDN can offer to next-generation networks, and describe a prototype solution based on extensions to OpenFlow. It will describe the architecture of Transport SDN and discuss various network virtualization techniques developed for different SDN applications. Concepts such as an open, virtual transport switch will be introduced, along with the different roles that GMPLS can play in Transport SDN. Finally, key issues and potential SDN-based solutions for multi-layer, multi-vendor networks will be presented.

.

# Contents

In the quest for expanding product capabilities, cable operators today are focused on delivering a broad array of services, products, and content to end customers anywhere, anytime, and on any device. The rapid growth and broad adoption of cloud based services and the delivery of rich content are 2 key drivers of data traffic within the backbone, but high-bandwidth cable backbone operations that enable greater service availability, network resiliency, and better application performance are also necessary to maximize the end-user experiences. While a certain amount of backbone traffic is expected to remain relatively constant and predictable, there is growing evidence of bursty traffic driven by big data applications, unpredictable one-time events, and the distributed nature of cloud networking and virtualization.

Typical WAN infrastructures are built using a dynamic, multi-purpose IP/router layer to deliver a variety of services. These router backbones historically were interconnected by leased wavelengths over the WAN, but in recent years, many MSO's have migrated some or all of their infrastructure to privately owned optical backbones. This enables MSO's to better control the delivery of bandwidth between routers, but these optical networks still largely deliver static capacity, relegating all dynamic bandwidth and traffic management to the IP/MPLS routers and the vendor-specific control plane protocols.

Given the rapid growth in cloud and content and the on-demand need for bandwidth, there is an increasing need for optimizing the WAN interconnectivity between sources and sinks for this data (e.g., data centers), and enabling rapid delivery of bandwidth whenever and wherever it is needed, without designing and over-provisioning the network for peak bandwidth. Furthermore, many cable operators are seeking network architectures that can support more rapid introduction of innovative services, and facilitate a level of programmability and automation that enables the network to dynamically deliver bandwidth services with specific service requirements to applications as needed. This not only enables operators to be more responsive in an increasingly competitive market, but also paves the way towards adding intelligence to the network and optimizing resource utilization in the multi-layer network. As part of that vision, Software Defined Networking (SDN) concepts can be employed not just for programming and controlling the packet switching/routing systems, but also for optical transport systems providing WAN bandwidth services. The extension of SDN concepts towards the transport network is called Transport SDN.

### The Evolving Optical Transport Network

The rapidly evolving optical transport landscape is changing the way many service providers are looking at their core networks. As the need to maximize fiber capacity increases, cable operators along with carriers are looking to new optical transmission technologies to enable more bandwidth over their WAN fiber infrastructure. Many are already deploying 100Gb coherent optical technologies as a first step, enabling multiple terabits of data transmission. Needing to scale fiber capacity even further, the industry is largely heading towards flexible optical super-channels (multi-carrier optical channels) with programmable modulation schemes to enable service providers to dynamically trade off distance with wavelength capacity.

The increased wavelength speeds are creating a new dynamic in transport networks. The evolution of wavelength speeds are aimed at maximizing fiber capacity, but many of the cable operators' backbone networks require bandwidth services substantially less than the wavelength speed. Even though the need to upgrade to 100Gb wavelengths exists today, the cable network has a need to deliver a broad range of bandwidth services, much of which continues to be dominated by 10Gb services. This raises the question of how best to maximize the utilization of new 100Gb optical wavelength resources and what technologies make most sense.

One of the innovations that has emerged to improve bandwidth management in the transport network is the convergence of WDM transmission along with Layer 1 OTN switching. OTN technology enables cable operators to channelize the 100Gb wavelengths into flexibly sized transparent tunnels and efficiently multiplex and groom different bandwidth service rates on to the optical wavelengths, maximizing the utilization of the 100Gb resources. This has the effect of decoupling the delivery of bandwidth services from the underlying optical wavelength technology, and when integrated with WDM transmission technology, provides cable operators with the ability to flexibly and rapidly deliver right-sized bandwidth services wherever and whenever they are needed. It is important to note that this is also operationally decoupled from optical planning and engineering, which is typically associated with the turn-up of new optical capacity. Since Layer 1 OTN switching functions in the digital domain, bandwidth services can be turned up (and down) near instantaneously.

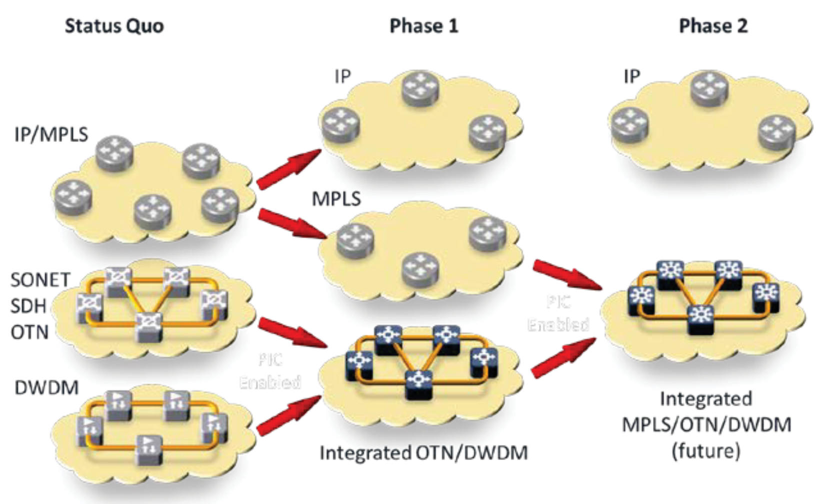


**Figure 1 Trends in network layer convergence**

This innovation in the optical transport domain has an important implication for bandwidth – instead of being a static resource that interconnects IP/MPLS routers, bandwidth can now be considered an elastic resource, capable of delivering capacity between points in the network whenever it is needed and for whatever duration makes sense. Furthermore, the incorporation of digital switching in the core can have significant implications on large packet flows which today are delivered through IP/MPLS routers and which could be more efficiently switched, transported, and even protected at the transport layer. This can help relieve the routers from performing simple transport functions.

However, one of the challenges facing operators today with multi-layer, multi-vendor networks is the lack of availability of a unified control plane. While control plane technologies today are widely deployed, they are isolated to specific layers and specific domains of the network, and implementations between vendors lack interoperability as well as consistency. This leads to a number of non-optimal situations, including:

- Non-optimized multi-layer networks: much of the networking functions, including simple transport and protection, are contained within just one layer, leading to a non-optimal total network solution. Transit traffic that could be switched or protected more cost effectively at the transport layer is instead left on the intermediate IP/MPLS routers, resulting in higher per-bit transport costs than alternatives.
- Lengthy operational cycles: in multi-layer and multi-domain networks, the non-existence of control plane interactions between network layers and domains often means serialized manual operations and coordination of resources between layers, leading to longer provisioning cycles.
- Non-optimal traffic engineering decisions: the lack of control plane interaction between layers means traffic engineering decisions are isolated to particular network layers (e.g., within the IP/MPLS layer) and as a result, are typically unable to take into consideration physical transport network characteristics such as bandwidth, latency and jitter.
- Excessive network over-provisioning: without taking advantage of the flexibility and adaptability the converged transport layer has to offer, networks are often designed for peak bandwidth, resulting in higher total network costs than is necessary.

## The Emergence of Transport SDN

Software Defined Networking (SDN) is an emerging architecture that is founded on the principal of separating control plane functions from the data forwarding plane, and enabling direct programmability of flows on packet forwarding hardware systems. One of SDN's objectives is to enable migration away from a model where the distributed control plane logic is tightly coupled with the vendor's hardware platform, and instead, create an alternative model where the networking intelligence is logically centralized, and the actual control over the flow of packets is shifted to the SDN control layer. Some of the benefits highlighted by the Open Networking Foundation (ONF) include [1]:

- Programmability – through protocols like OpenFlow, external controllers can directly program flow tables on packet systems, based on whatever routing criteria makes most sense for the network provider
- Agility – reconfigurability of the network can be done dynamically and based on existing or future network conditions
- Centralized Intelligence – the global perspective of the network can provide invaluable information that is necessary for performing any global optimization and for presenting a simplified abstraction of the network to higher level systems and applications
- Open Standards and Vendor-neutrality – industry efforts are underway by organizations such as the Open Networking Foundation (ONF) to create standards that will simplify interworking and establish common abstractions that are vendor neutral.

While much of the industry focus has been on enabling SDN for packet systems (and with a keen focus within the data-center), many of the same concepts of SDN are being logically extended towards the optical transport layers, which include packet, OTN, and WDM transmission technologies. The objectives of Transport SDN are multifold:

- Enable programmability of the flexible transport layer and leverage its increasing ability to switch and groom transport bandwidth over optical resources, as well as switch optical capacity
- Virtualize transport network resources and support a simple abstraction for provisioning bandwidth services
- Simplify, orchestrate, and automate provisioning operations within a multi-vendor, multi-layer, and multi-domain environment
- Enable the improvement of overall network resource utilization across multiple network layers
- Speed the delivery of new services and rapidly deliver the bandwidth on-demand to support these services, wherever and whenever it is needed

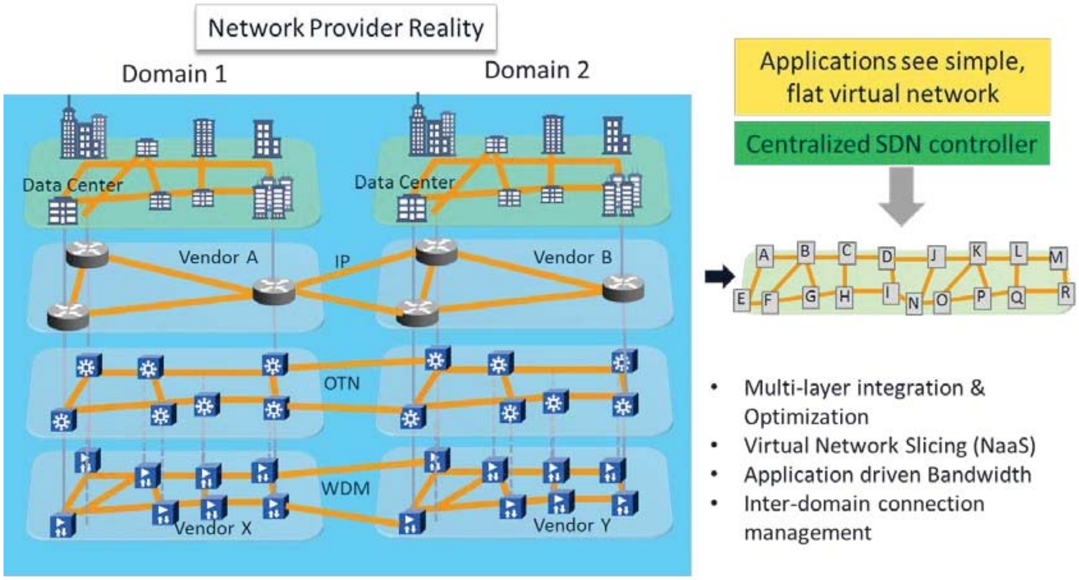


**Figure 2 Transport SDN enabled multi-layer, multi-vendor, multi-domain orchestration**

To achieve these objectives, an architectural framework is needed that logically centralizes topology information from multiple layers and which allows applications to see network resources across layers, vendors, and individual domains. Using this centralized topology data, more intelligent resource allocation decisions can be made, taking into consideration a broader set of service routing criteria (resource availability, latency, relative costs, power consumption, future bandwidth demands, etc.) than is currently available within vendor- and domain-specific control plane protocols.

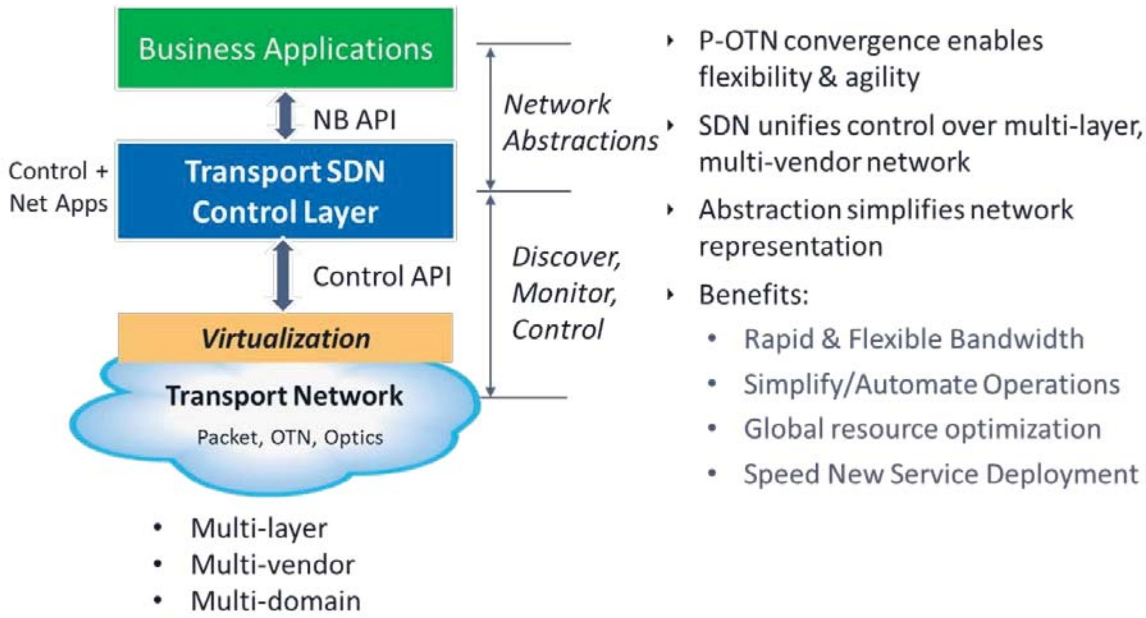


**Figure 3 Virtualization & abstraction play important roles in Transport SDN**

A key notion for enabling this architecture is the virtualization of certain network resources and support of abstractions of the network that simplify the bandwidth service provisioning function between on-network sites such as data centers. By virtualizing the underlying transport mechanism (whether it be MPLS, OTN, optical wavelengths, or other), applications would be able to program overlay bandwidth interconnections in a fashion that hides individual layer complexities. With this capability, and by leveraging the flexible nature of converged optical transport systems with integrated switching, true programmability of transport bandwidth flows over a mesh optical transport infrastructure can be achieved, thus enabling a more efficient bandwidth connectivity model that leverages the economics and capabilities the new converged packet-optical transport layer has to offer. A key part of this architectural framework is the concept of an Open Transport Switch.

## Open Transport Switch

The *Open Transport Switch (OTS)* [2] is an OpenFlow enabled, light weight virtual switch that manages physical resources associated with the optical transport layer. OTS abstracts network element (NE) resources which include ports and time slots/cross-connects into virtual resources which are exposed appropriately to an SDN Control Layer. Via a northbound API supported by this SDN Control Layer, applications can request end-to-end bandwidth services and aggregate packet traffic into optical trunks. This abstraction removes the need for control plane interactions at points of packet-optical transition creating a unified view of the network. This virtualized topology allows external applications to perform optimal path computation, provisioning and monitoring based on their QoS parameters. Though OTS manages optical transport resources, the concept can encompass the broader scope of core transport networks including MPLS and MPLS-TP.

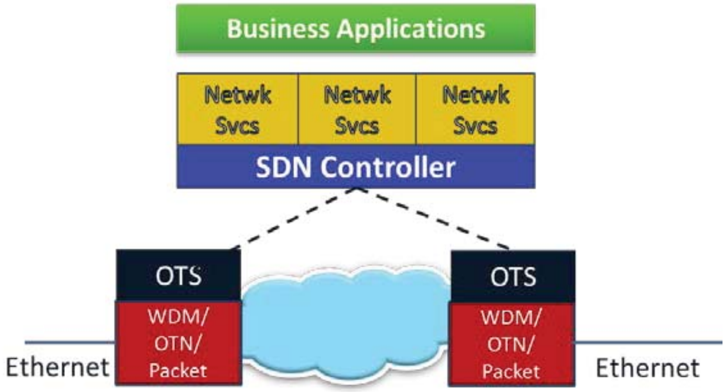


**Figure 4 OTS presents a virtual transport system abstraction to the Controller**

OTS is a key element in supporting network virtualization, as underlying physical resources and technologies can be virtualized and abstracted via its interfaces. Rather than directly managing technology-specific resources such as individual optical wavelengths, OTS can virtualize the optical capacity and support an abstraction that enables the SDN Control Layer and applications to view this as a flexible pool of available bandwidth, much like how storage and compute power is virtualized today in cloud networking. OTS provides support for functions central to bandwidth provisioning, including:

- Discovery of network topology and topology-related resources
- Monitoring of changes that impact topology, such as bandwidth availability changes and the availability of new resources (nodes, links, ports, etc.)
- Provisioning of bandwidth resources based on the provided service requirements, utilizing whatever underlying technology or capabilities are offered by the system, whether it be packet, OTN or WDM.
- Management and configuration of the OTS and its resources

As OTS is a virtual transport switch, it can provide different levels of abstraction relative to the physical network element. In one scenario, OTS can be used to represent the entire transport system NE, interconnected via optical capacity to other OTS instances that represent direct neighbors. OTS can also be used to represent a logical partition of a physical NE, enabling multi-tenancy on the transport system and the abstraction of a dedicated virtual transport switch with its own set of connection ports. These OTS instances can be interconnected with transport tunnels over the transport network, enabling support for virtual transport network topologies overlaying a physical transport network, sometimes referred to as Layer1 VPNs. Conversely, the OTS concept can also be extended to support the abstraction of a collection or domain of transport NE's, presenting the view of a single large logical (distributed) switch.

This virtualization of physical transport network resources allows resources to be pooled and then allocated, released, and re-allocated on-demand via software. OTS virtual switch instantiations and their connectivity can be flexibly configured based on the needs of the controller or applications and can enable bandwidth services to be programmatically provisioned.

## OTS Provisioning Modes

In order to satisfy the broad set of cable operator requirements around provisioning bandwidth services across the network, OTS accommodates different provisioning modes that integrate with existing control plane technologies to varying degrees. One of the reasons SDN advocates separation of the control plane from the forwarding plane is to enable external programmability of packet flows through packet systems, which historically have tightly coupled the control and forwarding planes into a single system.

Decoupling represents a significant paradigm shift for packet systems (switches, routers), which are inherently connectionless "store-and-forward" systems, switching or forwarding packets on to their "next hop" and letting the next system make its own decision on what to do next with the incoming packets. The development of embedded distributed routing protocols emerged as a necessity for connectionless systems just to make them function, since incoming packets need to know their next hop and which output port to head towards. With the emergence of SDN concepts, however, the paths of packet flows can be externally and centrally computed and then programmed into packet forwarding systems through standard protocols like OpenFlow, giving cable operators explicit control over how flows are routed through networks, rather than relying on vendor's control plane implementations.

In the connection-oriented world where optical transport systems reside, however, centralized control and programmability of circuits by external NMS/OSS systems have been in practice since the emergence of such technology. In the past decade this circuit-oriented model has been complemented by distributed control plane technologies like GMPLS, enabling an incremental level of network intelligence that did not exist before. GMPLS technologies enable multiple networking and automation features for transport networks that many operators utilize today, including:

- automated topology discovery
- accurate tracking and reporting of bandwidth resources
- integrated dynamic route computation for end-to-end circuits
- traffic engineering with support for explicit routing
- robust signaling for establishing end-to-end data paths
- high resiliency from multiple network failures through dynamic restoration of circuits

Transport SDN extends SDN's openness and programmability concepts to the transport domain, and has the potential to be more easily implemented because of existing logically centralized models for directly programming connections. This does not imply vendors will simply reuse/repackage legacy management systems, but rather the fundamental notion of a centralized approach already exists in the transport world, providing a potential time to market advantage as it is less of a paradigm shift.

For facilitating different abstractions and catering to different operator operational models, OTS offers multiple provisioning modes. One mode of operation is called **explicit** or **direct mode**. In this operational mode, the SDN Control Layer maintains knowledge of the full transport layer topology, along with a path computation function that leverages whatever metrics makes business sense. Once an appropriate route is computed for the bandwidth service, the SDN Control Layer programs each of the OTS instances along the path in a hop-by-hop manner. When the topology information from multiple domains is logically centralized, the SDN Control Layer can easily determine a path across the multi-domain network.

Another important operational mode OTS supports is the **implicit mode**. This mode leverages a pre-existing control plane within the transport network, such as GMPLS. Many operators today have large optical network deployments that utilize GMPLS to varying degrees and have an inherent interest in preserving this level of intelligence, but would also like the option of evolving towards an SDN-like model.

In this mode, the SDN Control Layer has a view of just the edge nodes of the domain (where bandwidth services can originate or terminate), and the transport network domain is abstracted as a single network fabric. In this abstraction, the route computation is relegated to the embedded control plane, based on provided service parameters and SLA requirements. When bandwidth service is required, the SDN Control Layer sends a request to the edge OTS, specifying the remote service endpoint. The lion's share of the work to compute a path through the network fabric and set up the optical flow is handled by GMPLS. In a multi-domain scenario, the SDN Control Layer can issue provisioning requests to each domain, and in this fashion, can stitch together multiple segments to create an end-to-end transport flow.

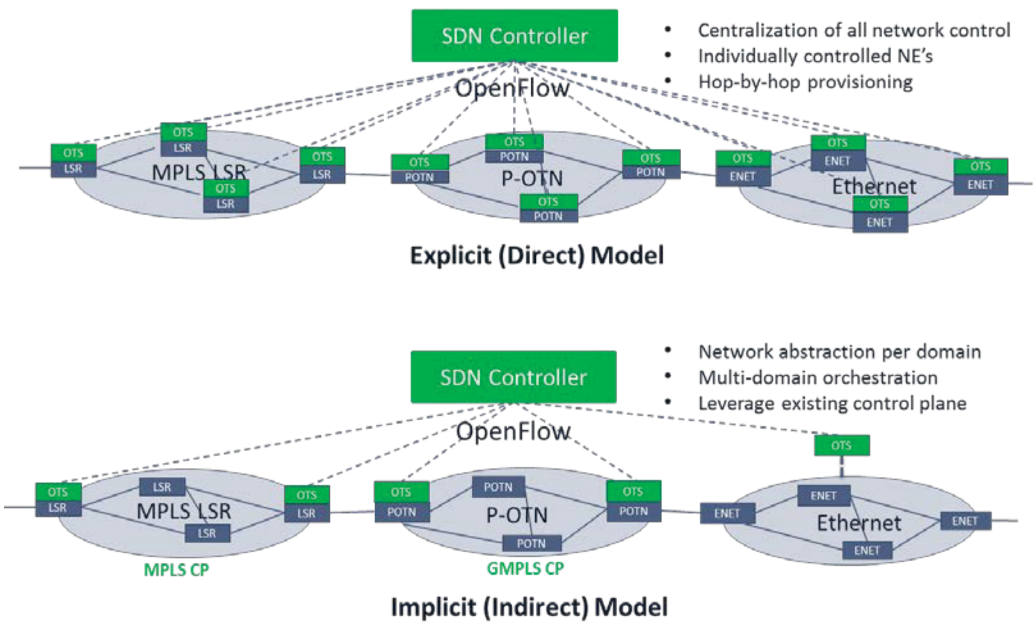


**Figure 5 OTS supports multiple provisioning modes & facilitates compatibility with pre-existing control planes in transport domain**

Various **hybrid modes** combining aspects of explicit and implicit operational modes are also feasible and enable further adaptation to operator provisioning processes. One such hybrid model centralizes the path computation function where a global view of the network resources resides, but activation of the bandwidth is then pushed down into the network as an explicit or directed path. This leverages the robust path set-up capabilities of GMPLS, rather than relying on hop-by-hop requests sent from the controller down to each network element.

## Convergence and Multi-Layer Networking

The quest to simplify and reduce total network costs is leading to the convergence of networking layers. Today, we already have commercial WDM transport systems that efficiently integrate Layer 1 (OTN) switching functions, thus enabling operator backbones to switch, multiplex, and groom flexibly sized ODUflex circuits. These circuits provide transparent tunnels for higher layer networking layers, and can be efficiently groomed to maximize the fill of 100Gb wavelength resources. In an environment where 10G rate services and interfaces still dominate, OTN serves an important role in enabling flexible grooming of bandwidth services into higher-rate optical carriers. When integrated OTN switching comes with very little incremental cost above and beyond the WDM transport itself, it can provide an invaluable networking capability that can help reduce the cost of the total network solution [3]. Converged WDM/OTN systems enable:

- efficient grooming of multiple service rates (e.g., 10, 40 & 100Gb router interfaces) onto 100Gb wavelengths
- rapid turn-up and turn-down of transport bandwidth services through deterministic digital switching and the decoupling of bandwidth service delivery from wavelength engineering
- new cost-effective protection & restoration alternatives to MPLS FRR, such as sub-50ms Shared Mesh Protection
- reduction of the “aggregation tax” associated with hierarchical packet topologies by enabling higher degree of meshiness and offloading larger transit traffic flows from intermediate routers

Network convergence is extending beyond just Layers 0 and 1 – already the industry is starting to see convergence of packet and optical technologies, primarily in the metro environment. This trend is gradually migrating to the core, where packet functions and switching (up to Layer 2.5) are being integrated along with Layers 0 and 1.

Today, it is not uncommon to see multi-layer networks comprised of an IP-MPLS router layer and an optical transport layer that either uses ROADMs to switch wavelengths with statically multiplexed circuits or which uses digital OTN switching to provide any-to-any bandwidth service delivery. Despite advances in network architecture, the management and operations of multi-layer networks are still very much done layer by layer. Functions including planning, engineering, turn-up, and management within multi-layer networks are still achieved today by separate organizations with little automation. The lack of interactions between network layers often necessitates manual processes for performing basic functions such as adding new IP link capacity. This lack of an integrated view of the multi-layer network also means little opportunity to perform global optimization.

While there have been attempts in the industry to enable intelligent cross-layer and inter-domain control planes in a multi-vendor environment, these efforts have for the most part been unsuccessful commercially. Transport SDN introduces an alternative means for enabling multi-layer orchestration and optimization by centralizing the control planes and allowing higher level applications to directly program services across different network layers, across domains, in a vendor-neutral fashion.

### Enabling Transport SDN for Multi-layer Networks

To support this notion of a centralized network control plane, the SDN Control Layer needs to incorporate topological information from each network layer and gain an understanding of the relationship and dependencies between the layers. Additionally, it requires a southbound programmable control interface to provision data path connectivity at the appropriate locations and at each network layer. Without an understanding of the topological relationship between layers, it is difficult for any path computation algorithm to compute a globally optimal solution. Once this topological data is centralized, multiple link metrics can be applied as appropriate, providing the path computation application multiple different route optimization criteria. In addition to hop-count and administrative costs, metrics such as link economic cost, nodal transit costs, and latency can be incorporated into the model.

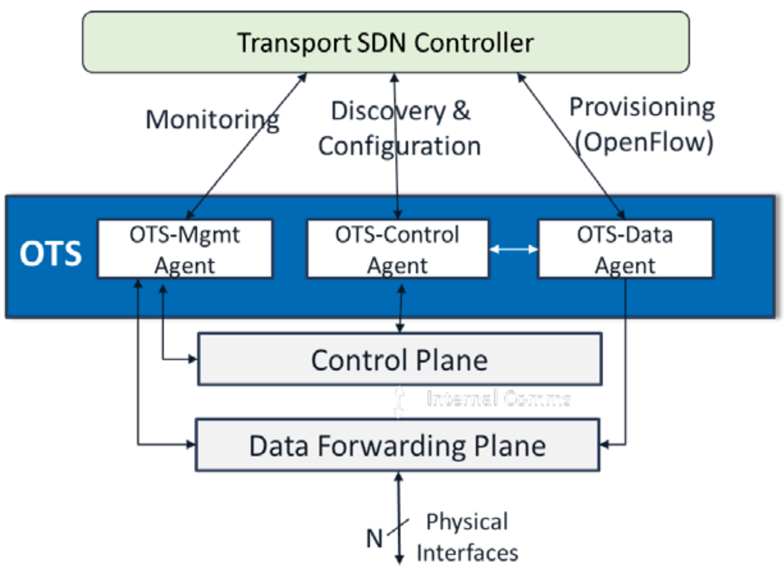


**Figure 6 High level OTS functional architecture**

For the converged optical transport layer, the OTS plays a key role in enabling Transport SDN. OTS, as a thin software agent, can run natively on the NE and supports an OpenFlow based interface for integration with the SDN Control Layer. It is comprised of 3 main components:

- ***Data Plane Agent***: This agent is responsible for programming the NE data path to create/update/release circuits. The data path entities could be cross-connects or MPLS labels or VLAN tags and are activated using OpenFlow wire protocol.

The agent also supports other aspects of the OpenFlow wire protocol, such as retrieval of OpenFlow resources and reporting status updates on objects.

- ***Control Agent***: this agent is responsible for topology discovery and management. It has responsibility for informing the SDN Control Layer of current topology (links, ports) as well as ongoing topological changes. The Control Agent monitors and propagates topology-related notifications and alarms to the SDN Control Layer to ensure its Traffic Engineering (TE) database and state are kept consistent with the network element's view.
- ***Management Agent***: the management agent is responsible for configuring and managing the OTS instantiation, which includes adding/removing ports from the OTS and managing the OTS lifecycle. This agent is intended to support an OpenFlow interface that is analogous to the OFConfig interface used for packet systems.

## Optimizing Multi-Layer Networks via Transport SDN

Through the interfaces to OTN, it is easy to see how a multi-layer SDN controller can extract the transport layer topology (which may be either physical or virtual, depending on how OTS is instantiated) and use this for building a multi-layer topology database. Other mechanisms may be used to import or deduce the topology from other network layers.

The linkage between the layers, however, is something that is not easily automated with today's technology. Well known auto-discovery protocols do exist for protocols like Ethernet (e.g., LLDP), and can be leveraged to auto-discover Ethernet-based inter-layer connectivity, and converged transport systems are beginning to incorporate packet-processing capabilities into their optical tributaries, but in transport systems that are fully transparent to Ethernet (e.g., traditional transponder or muxponder), manual configuration of the connectivity relationship between switch/routers and the transport systems is often required.

Once the SDN Control Layer builds up a view of the multi-layer topology, and a global view of network resources is attained, applications can be constructed to optimize traffic flows over the network and automate the provisioning of the multi-layer network. One use-case scenario is facilitating efficient transport of big data. In many networks, flows associated with particular activities can create large demands for bandwidth, sometimes for fairly long durations. Events such as Virtual Machine migrations associated with cloud computing, internal transport of large data sets, and data/storage backups can generate a substantial amount of traffic over the core network at varying times. These data flows, sometimes referred to as "elephant flows" possess traffic characteristics that differ in comparison with the many smaller transient flows more commonly associated with end-user interactions with the Internet such as web browsing, live streaming, etc.

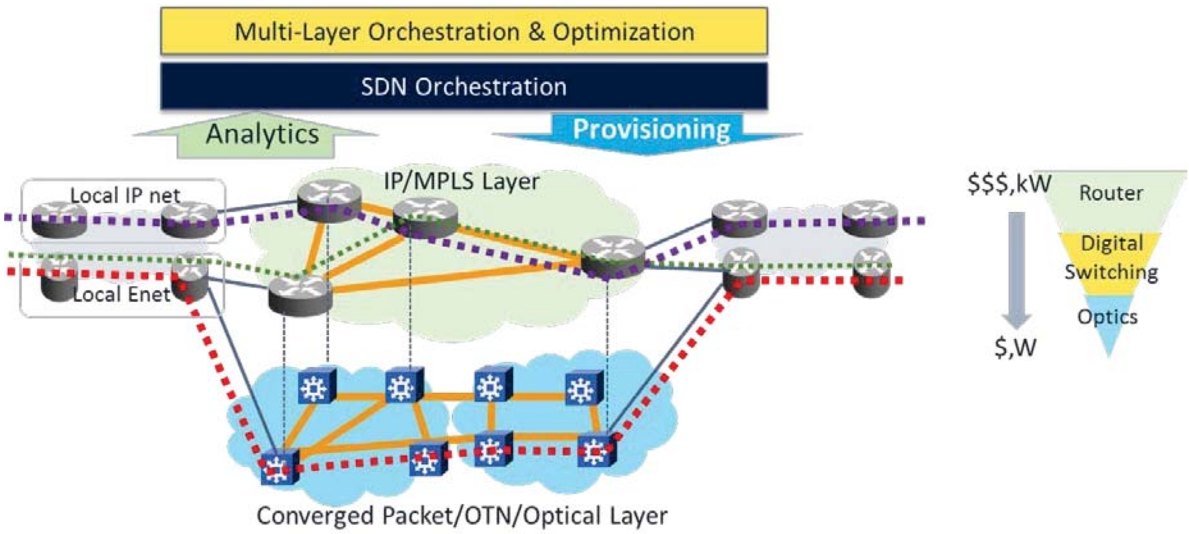


**Figure 7 Transport SDN facilitates orchestration of flows across multiple layers for optimal resource utilization**

One common approach to architecting a network to support this traffic is to rely predominantly upon the IP/MPLS packet layer to handle all traffic management functions, regardless of the flows' characteristics. While simple from an operations perspective, it can lead to over-dimensioning the network and a non-optimal use of resources, particularly when the economics of transporting a bit through a core router is higher than switching that same bit through a converged transport system. For elephant flows, it is often more economical to offload that traffic from the packet layer and all of the intermediate routers and instead deliver it via the transport layer. Leveraging the flexibility of the new converged transport layer to dynamically provision any-to-any bandwidth and directly switch "circuit-like" traffic flows helps off-load the IP/MPLS layer from performing transport functions. Identifying an optimal solution that leverages the available resources at multiple network layers is the essence of multi-layer traffic engineering, and a key application for Transport SDN.

In order to operationalize this type of multi-layer optimization, the Transport SDN solution requires some capabilities, including:

- **Multi-layer Traffic Engineering** – an application that can build a logically centralized multi-layer topology database and route a flow through it, taking into consideration various routing metrics can help operators leverage the best capabilities of each network layer. These may include latency, power consumption, actual cost/bit, etc.
- **Network Analytics** – an application that can monitor existing traffic flows and characterize them is needed to identify candidates for multi-layer path re-engineering.
- **Multi-layer Control and orchestration** – in order to activate the engineered data path with minimal human intervention, the SDN Control Layer needs to coordinate and provision the relevant network layer resources through supported

(and ideally, standardized) Application Programming Interfaces (APIs). ONF defines standard protocols such as OpenFlow for programming flows through OpenFlow switches and is currently defining OpenFlow interfaces for optical transport systems. In router networks where OpenFlow is not supported, the SDN Control Layer could utilize vendor-specific, non-OpenFlow based interfaces to ideally achieve similar functionality.

## Conclusions

The emergence of converged optical transport technologies such as P-OTN and the development of new approaches to controlling network traffic leveraging SDN concepts are together creating new ways for cable operators to build multi-layer core networks that deliver bandwidth services not only more flexibly, but also more efficiently. The integration of OTN and packet switching technologies into optical transport platforms is resulting in a highly agile and multi-functional network layer that creates opportunities to offload transit traffic from the higher network layers (e.g., IP/MPLS) and thereby enhance the total cost of the networking. P-OTN technology provides cost-effective and efficient grooming of multi-rate traffic streams into optical super-channels, minimizing stranded capacity and enabling bandwidth services to become decoupled from the underlying optical transmission resources.

With integrated switching capabilities, P-OTN systems today are enabling optical mesh networks to be created, flexibly sized transparent circuits in increments of 1Gbps, and fostering new innovations such as Shared Mesh Protection, which enables bandwidth-efficient protection schemes in the optical transport layer. Collectively, this provides a new set of networking tools to help cable operators flexibly, dynamically, and cost-effectively transport higher network layer services, as well as protect them from network and equipment failures, resulting in increased network resiliency and an enhanced end-user experience.

In addition to networking advances within the transport layer, and the network virtualization capabilities P-OTN is introducing to operator backbone networks, the operationalization of the evolving multi-layer network must also be addressed. Distributed control plane technologies such as GMPLS have helped in advancing transport networking, but predominantly within single-vendor domains. Transport SDN, as we've described in this paper, is an emerging architectural approach that addresses the challenge of multi-vendor, multi-layer, multi-domain networks, leveraging key SDN principles such as the logical centralization of intelligence and letting applications directly program bandwidth services in the network, wherever and whenever required. A key concept discussed in this paper is the virtualization of network resources which enables operators to pool transport network resources together and allocate right-sized bandwidth on demand in an efficient manner. With the centralized knowledge of network resources across domains and layers, Transport SDN provides the fundamental building blocks needed for optimizing network resource usage and tailoring the traffic flows to take advantage of the best each network layer has to offer.

# Bibliography

[1] Open Networking Foundation, https://www.opennetworking.org/sdn-resources/sdn-definition

[2] A. Sadasivarao et al, '*Bursting Data Between Data Centers: Case for Transport SDN*', 21st Annual Symposium on High-Performance Interconnects (HOTI 2013), to be published.

[3] A. Deore *et al*, *'Total Cost of Ownership of WDM and Switching Architectures for Next-Generation 100Gb/s Networks'*, IEEE Communications Magazine, Vol. 50, No.11 (Feature Topic Issue on High-Speed Optical Transport), November 2012

# Abbreviations and Acronyms

| Acronym | Description |
|---|---|
| API | Application Programming Interface |
| FRR | Fast ReRoute |
| Gb | Gigabit |
| GMPLS | Generalized Multi-Protocol Label Switching |
| IP | Internet Protocol |
| MSO | Multiple System Operator |
| LLDP | Link Layer Discovery Protocol |
| MPLS | Multi-Protocol Label Switching |
| MPLS-TP | MPLS Transport Profile |
| NE | Network Element |
| NMS | Network Management System |
| ODU | Optical Data Unit |
| ODUflex | Flexible rate Optical Data Unit |
| OFConfig | OpenFlow Configuration and management protocol |
| ONF | Open Networking Foundation |
| OSS | Operations Support System |
| OTN | Optical Transport Network |
| OTS | Open Transport Switch |
| P-OTN | Packet Optical Transport Network |
| ROADM | Reconfigurable Optical Add Drop Multiplexor |
| SDN | Software Defined Networking |
| SLA | Service Level Agreement |
| TE | Traffic Engineering |
| VLAN | Virtual Local Area Network |
| VPN | Virtual Private Network |
| WAN | Wide Area Network |
| WDM | Wavelength Division Multiplexing |