Integrated OTN Switching Virtualizes Optical Networks

By Andrew Schmitt
Principal Analyst, Optical
Infonetics Research, Inc.

INFONETICS RESEARCH

OTN AND THE OPTICAL REBOOT

A big upgrade of any system requires a full reboot; the changes are so extensive that a rebuild from the ground up is unavoidable. Optical networks are no different, and such a change is imminent. The broader equipment supply chain and investment community are now aware of the transformational event that will take place in 2013–2014—what we first dubbed the "optical reboot."

The introduction of coherent 100Gbps (100 Gigabits per second) optical transport is a key catalyst; the technology offers massive performance gains over incumbent technologies to extract more fiber capacity, but requires that carriers undertake greenfield builds to extract maximum return. Beyond achieving cost per bit gains with 100Gbps, carriers are seizing this opportunity to roll out new architectures that will further improve the total cost of ownership of their networks. The most important of these is a more sophisticated and efficient architecture for switching and managing these optical networks, which will use 100Gbps wavelengths efficiently and allow carriers to meet future operational cost targets.

The ITU G.709 standard for OTN (optical transport network) is the protocol of choice; it will bring the efficiency, reliability, and predictability of a transport-based approach and meet the infrastructure requirements of packet networks. OTN offers many features that will help facilitate the transparent transport of tomorrow's growing packet traffic and support services such as private line Ethernet or wavelength services requiring hard service level agreements. OTN is just as important to the optical reboot as 100Gbps and coherent optics.

Some in the industry have talked about eliminating the OTN switching layer and deploying a pure IP or MPLS solution with colored optics feeding into a simple WDM line system. These ideas mimic IPoDWDM architectures proposed 10 years ago that were never deployed in material volume. Our discussions with service providers indicate that the transport layer—specifically with OTN switching—is and will continue to be a fundamental part of future network architecture.

Service providers may desire to deploy a hybrid MPLS and OTN solution in the future, but today all 100Gbps WDM networks are being built using OTN transport, and as service providers build out these networks, most plan to employ OTN switching, which eventually will underpin the efficient convergence of packet and optical transport functionality into a single network layer.

VIRTUALIZATION FOR OPTICAL NETWORKS

Virtualization is the practice of abstracting a pool of common resources for use by multiple services—virtualization has revolutionized computing in enterprise and large datacenter applications. Virtual machines (VMs) in the data center allowed computing resources (hardware) to be decoupled from applications/services, improving hardware utilization and efficiency to save capex. VMs also reduce opex by bringing centralized management of these hardware resources and services. Services with different priorities cab be quickly shifted and cloned on top of a homogenous hardware resource pool.

These concepts have been so successful inside the data center that carriers would like to extend them into the network, allowing them to deploy any service into a pool of optical resources, where those optical resources can be quickly increased and decreased on demand. We need a transport protocol that is service-independent to realize this vision of a virtualized optical network—something able to carry everything from legacy SONET/SDH to Ethernet, MPLS, IP, and other protocols such as Fibre Channel.

There must also be a method for partitioning the massive bandwidth of 100Gbps channels or 1Tbps (1 Terabit per second) superchannels to efficiently pack a mix of protocols and speeds from multiple customers.

OTN transport provides an elegant solution for carrying and managing traffic in such a virtualized optical network by providing a standardized digital wrapper that can carry a wide range of services transparently across the network. **OTN switching** adds the key features for partitioning and grooming bandwidth among services with varying capacity, latency, and reliability requirements, and most importantly enables the multiplexing of many low-medium data rate protocols into efficiently filled 100Gbps wavelengths. Put together they form the optical transport technology that is the bedrock for future networks.

OTN TRANSPORT

There is some confusion about OTN, as the term is often generically used to describe three functions: OTN transport, OTN multiplexing, and OTN switching. They are not interchangeable; for example switching is a superset of transport functions, and it is worth examining the differences.

OTN transport has been in use for a decade, even earlier if one considers the early implementations of G.975 framing in SLTE applications in the 1990s. G.709 OTN was originally defined as a point-to-point protocol, designed to provide a protocol-independent wrapper of client data. The objective was to use a single homogenous protocol to wrap/containerize various clients, providing 100% transparent transport, something SONET/SDH was incapable of for Gigabit-speed services such as Ethernet, Fibre Channel, and wavelength services. Transparency is an essential element of transport networks; it means any protocol can be delivered with guaranteed bit rate and no change to the payload whatsoever. OTN is prized as a way to carry someone else's bits without modification.

Wrapping the source signal in OTN also allows operators to add extra bits—for operations, administration, management, and provisioning (OAM&P)—to the container and cleanly remove them at the destination. G.709 transport added these needed management features for today's networks where signals might transit multiple operator networks during transport. OTN transport provides a perfect service-independent, network-to-network interface for transiting the multiple transport domains within a large service provider or among multiple providers. As more operators adopt OTN, particularly the switching features, it will become the common currency for transport handoffs between different carriers.

Exhibit 1: OTN Container Payload Rates and Supported Clients

OTN Container	Payload Rate	Supported Clients
ODU-0	1.25Gbps	STM-1/4, Gigabit Ethernet, Fibre Channel, others
ODU-1	2.5Gbps	STM-16, 2 Gigabit Fibre Channel
ODU-2/ODU-2e	10Gbps	STM-64, 10 Gigabit Ethernet
ODU-3	40Gbps	STM-256, 40 Gigabit Ethernet
ODU-4	100Gbps	100 Gigabit Ethernet
ODU-FLEX	N x 1.25Gbps	Sub-rate Ethernet, or constant bit rate clients like EPON, GPON, CPRI, etc.

MUXPONDERS: THE GOOD AND THE BAD

Not long after G.709 gained traction, companies extended the protocol to wrap not just a single client, but multiple clients, allowing a single 10Gbps OTN container (OTU2) to carry multiple sub-containers with a mix of clients such as OC-48, Gigabit Ethernet, and Fibre Channel. This is **OTN multiplexing**, and it allowed OTN hardware to move beyond a point-to-point transponder role to multiplex data between two points (hence the term "muxponder").

Muxponders provided a vital function in the last decade as 10Gbps costs dropped and these wavelengths could be repurposed to carry the majority of clients between 1Gbps and 2.5Gbps. Our research indicates 85% of all deployed 40Gbps links use 4x10Gbps muxponders. But muxponders are a static solution; to make any changes to incoming and outgoing OTN containers, a human must change the fiber connections. To interface multiple WDM links, more muxponders are placed back to back and more complex cross cabling are required.

Exhibit 2: Muxponders Are a Static Solution

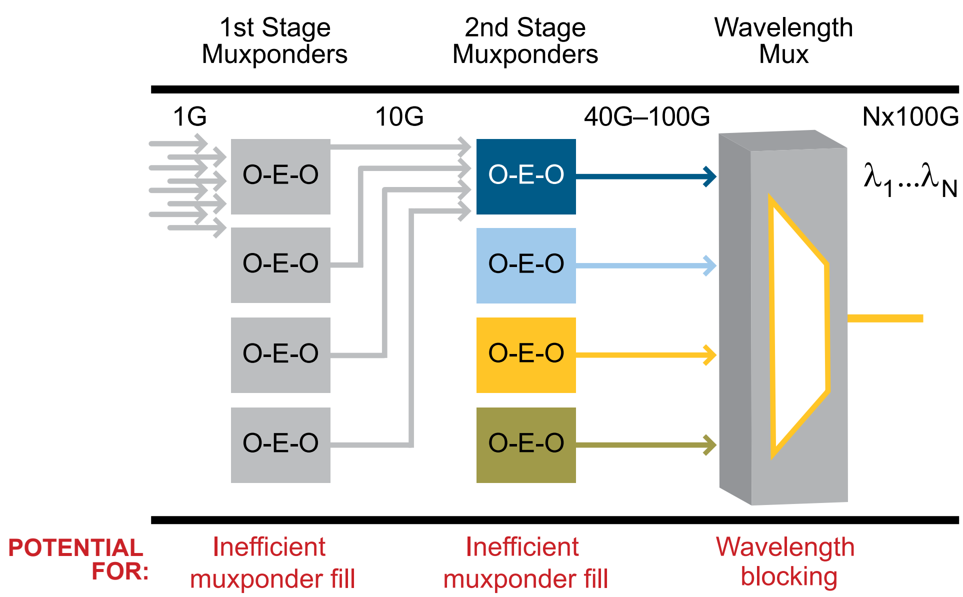

The shortcomings of the muxponder approach become visible as more are deployed. What was once a cost-saving solution to squeeze maximum efficiency out of a point-to-point link becomes unwieldy in scale, and an increasingly inefficient way to transport services as WDM data rates scale from 10Gbps to 40Gbps and 100Gbps and beyond. A node with multiple muxponders must be manually patched, adding capex and opex, and the architecture can't handle a dynamic mix of <=10Gbps services with 10Gbps, 40Gbps, and 100Gbps wavelengths in concurrent use.

Worse, muxponders provide only point-to-point transport capability, and even when paired with ROADMs, they do not efficiently allow the grooming of services within wavelengths or between wavelengths. Optical switching via ROADM only allows wavelengths to be switched and cannot add, drop, switch, or even monitor the multiple clients that might be sharing the capacity of a 100Gbps wavelength.

Muxponders also lock operators into the long legacy of manual provisioning, an error prone process that cannot adapt to the efficiencies required of future meshed networks. As optical transport networks become more meshed in an increasingly connected world, large deployments of static, inflexible muxponders perpetuate an architecture that wastes capacity, is difficult to manage, and cannot evolve into what carriers want—a virtualized mesh based optical network.

INTEGRATED OTN SWITCHING TO THE RESCUE

Muxponders and G.709 provided the technology foundation for **OTN switching** by defining a protocol that allowed multiple clients to be transparently bundled into uniform containers and sent on a single wavelength. OTN switches represent a quantum leap in architecture by a common electrical switch fabric—**OTN switching**—for hundreds or even thousands of wavelengths and clients to cross-connect at a particular node.

The first advantage of evolving from transponders and muxponders to OTN switching platform is the decoupling of clients from the wavelength transport interface. Rather than having a bundle of 4 or 10 client ports hard-wired to a single WDM line interface, each of the clients can be individually routed to any WDM interface or client port on the system. Even the few operators with no intention of using network-wide OTN switching find this capability valuable, as clients can be remotely provisioned without sending engineers to patch fiber cables.

Exhibit 3: With OTN Switching, Clients Can Be Individually and Remotely Routed to Any WDM Interface or Client Port

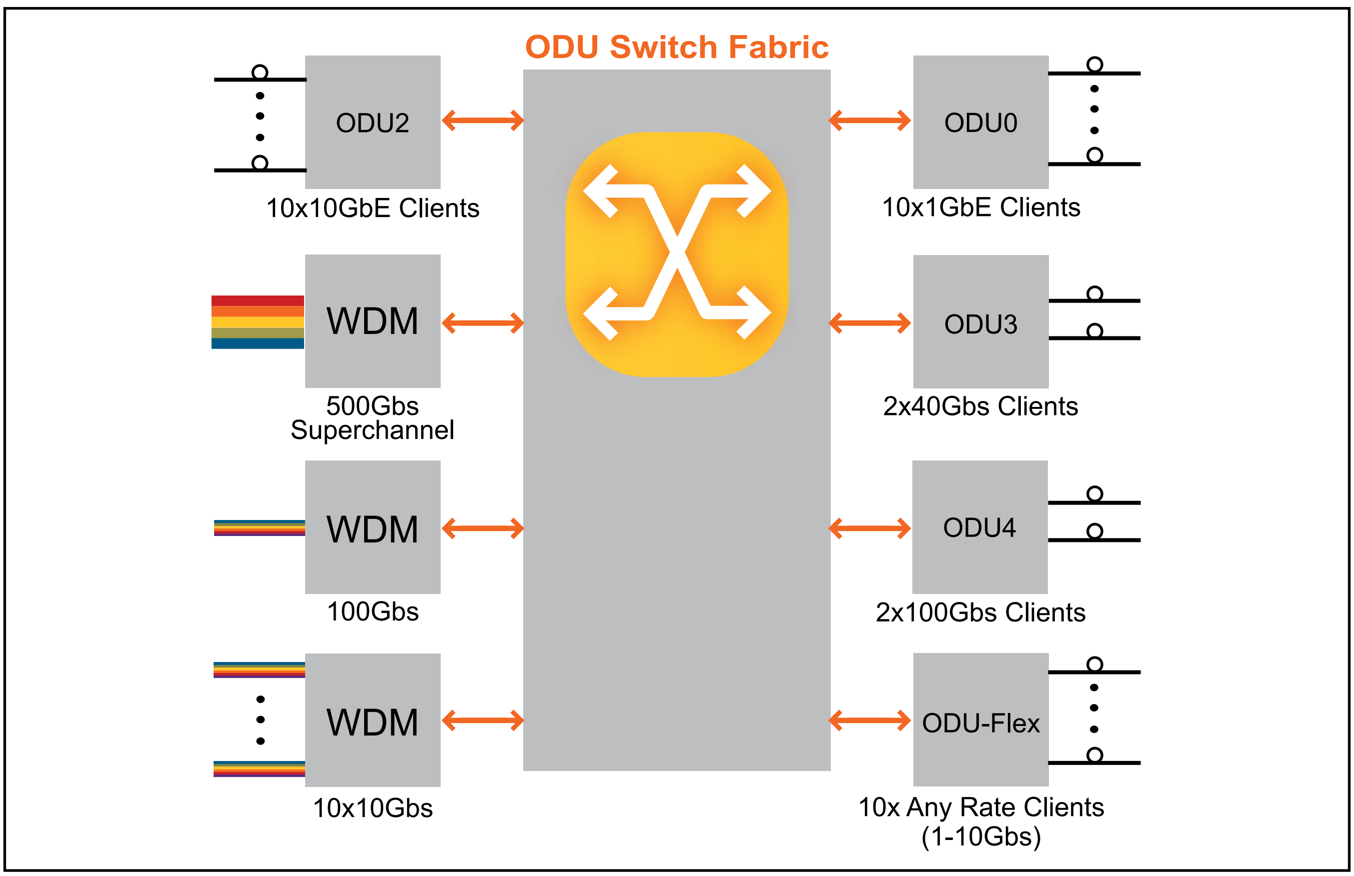

"Networks built with integrated OTN switching bring the economic and operational benefits of virtualization to optical networking."

Decoupling the clients from the transport interface also increases network efficiency. We live in a world where the dominant service interfaces sold by carriers are 10Gbps and lower and will be for many years, because the economics of 10Gbps clients remain tough to beat. 100Gbps is just too big a pipe for all but the largest enterprises and data centers. OTN switching allows multiple 10Gbps signals from disparate locations to be efficiently packed into these higher speed wavelengths, enabling fiber capacity scaling while simultaneously maximizing bandwidth utilization and efficiency. This is important for 100Gbps networks, but becomes vital if flexible-coherent schemes are used to implement superchannels at even higher data rates such as 200Gbps, 500Gbps, and 1Tbps.

The scaling issues of muxponders become apparent if, for example, 100 10Gbps interfaces needed to be hard-wired into a 1Tbps superchannel. The reality is this superchannel—comprising a hierarchy of sub-containers, and using a width of multiple optical channels—would carry a mix of everything from Gigabit Ethernet and Fibre Channel (ODU0), OC-48 (ODU1), 10Gbps (ODU2), perhaps a stray 40Gbps client (ODU3), and several 100Gbps trunk lines (ODU4). This 1Tbps of traffic would originate at multiple client and line WDM interfaces—an integrated OTN switch fabric on the node with the 1Tbps superchannel is the only way to realistically bundle 1Tbps worth of smaller ODU containers. Many of these services may have different origination and termination points, and thus if they were assigned to muxponders there would be no way to efficiently groom them onto a common superchannel, thereby requiring overbuilding of more inefficiently filled WDM capacity.

BENEFITS OF VIRTUALIZED OPTICAL NETWORKS

Networks built with integrated OTN switching bring the economic and operational benefits of virtualization to optical networking. These benefits are found in four areas: capacity, service velocity, provisioning, and failure restoration.

Capacity: As illustrated earlier, OTN switching decouples the clients from the WDM line interfaces, allowing greater network efficiency by ensuring that the more costly WDM links are running as hot as possible and that no stranded bandwidth remains. A network of OTN switches takes this concept further, allowing traffic to be aggregated at intermediate nodes and directed towards underutilized routes. Though the nonstop flight from Boston to San Francisco may be full, you might find unused capacity on the multitude of flights with connections. OTN switching allows traffic to move through intermediate network nodes transparently and puts unused capacity to work.

Service velocity: Virtualization allows new services to be quickly added in the data center, and OTN switching does the same for transport networks. The presence of switches at each node makes it easy for new clients to be attached without concern to which muxponder card is being addressed. Service changes can be processed just as easily and released capacity returned to the pool of virtualized optical resources.

Provisioning: There may be special requirements for clients requesting capacity from a virtualized optical network. As new services are turned up latency, protection, and policies must be considered when the connection path is computed. Integrated OTN switching allows a mesh based approach to provisioning, allowing multiple clients sharing the same virtual optical network to take paths that meet each customer's specific requirements. Financial traffic can be provisioned on the lowest latency route, for a premium. Intra data center traffic can be provisioned with deterministic latency, so that multiple available paths fall within specified constraints. Government or military traffic can be constrained to avoid certain areas and identical paths through the network. The presence of OTN switching throughout the network creates a very meshy resource, which gives the control plane more options to meet specific customer requirements.

Restoration: Network connections inevitably fail, and the same control plane that originally provisioned the service must re-route traffic upon the loss of a connection. Some or all of the same provisioning requirements must be met again, but with fewer transport resources available this time. Again, having a meshy network with ubiquitous OTN switching on all available resources will allow the most optimal solution to be generated. OTN switches can be rapidly reconfigured, much faster than ROADMs, allowing restoration to take place as quickly as possible.

Exhibit 4: Integrated OTN Switching Benefits: Capacity, Service Velocity, Provisioning, Failure Restoration

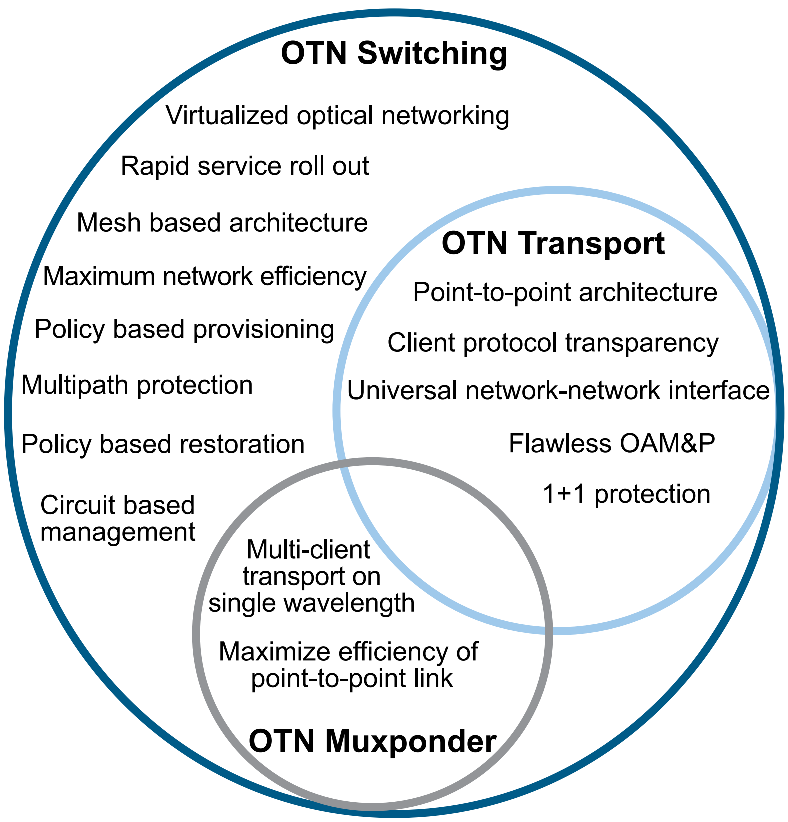

"Carriers large and small that want to build efficient transport networks are adopting OTN switching."

WHO WANTS OTN SWITCHING

Most carriers want OTN switching, and these carriers represent the vast majority of global service provider capex.

In April 2012, Infonetics concluded a survey of 21 global service providers, together representing 34% of global operator capex. We asked these operators whether they use or plan to use electrical OTN switching or use OTN strictly as a transport protocol; 76% plan to deploy OTN switching or have already. These respondents account for a whopping 90% of respondent capex, indicating the larger carriers are embracing OTN switching. Conversations with these carriers cite a number of reasons OTN is well suited to their networks.

Exhibit 5: Carriers That Plan to Use OTN Switching Typically Have Greater Capex Than Those That Do Not

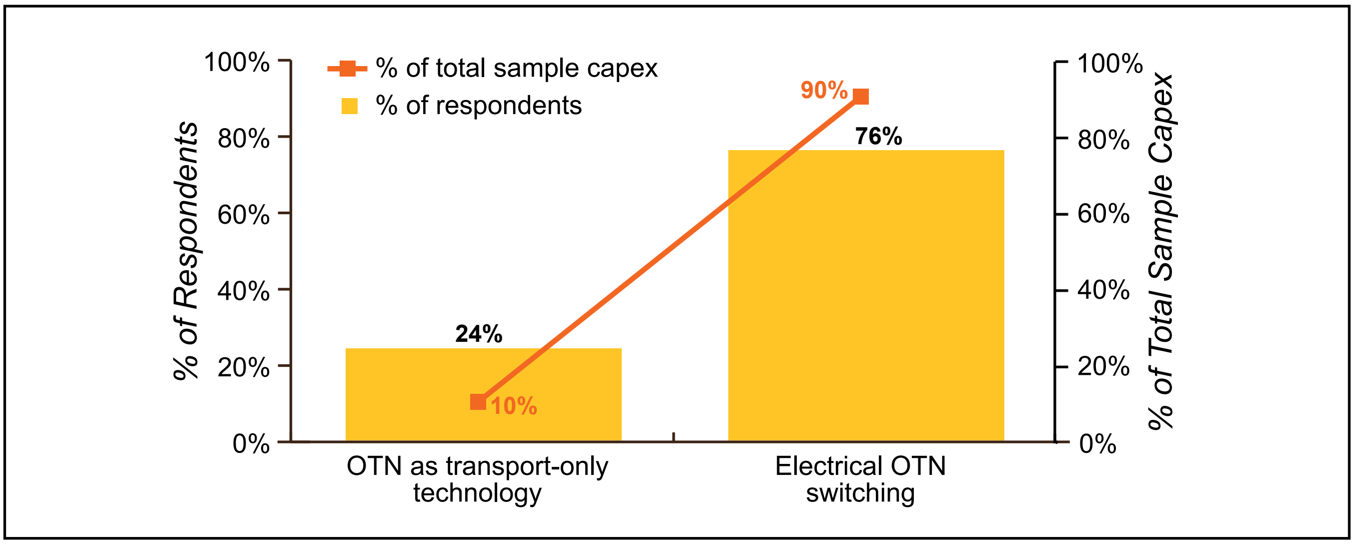

Source: Infonetics Research, *OTN Deployment Strategies: Global Service Provider Survey*, April 2012

One common need that emerged in conversations with the larger carrier respondents is a requirement to unify their transport infrastructure and operate multiple networks on single common transport infrastructure. Wireless, broadband, enterprise data services, and leased line wavelengths can share wavelengths via OTN transport. OTN switching coupled with an intelligent control plane allows each service to be routed and re-routed according to required constraints.

Other carriers cited a desire to move to a more advanced control plane, with OTN switches providing the homogenous transport infrastructure and mesh intelligence that current disjointed SONET/SDH networks cannot. The transition to 100Gbps wavelengths is the catalyst for this transition, and mesh-based OTN switching is the best way to fill these fat pipes.

Existing 10Gbps network architectures worked well for a while, but all the low-hanging fruit has been picked—carriers recognize the need for a break from the past, a "reboot." They are seeking maximum efficiency in their optical networks as they struggle to meet traffic loads increasing 30% to 50% per year on fixed capex budgets. Carriers large and small that want to build efficient transport networks are adopting OTN switching.

"If DWDM transport and switching functions can be combined without compromise at the right cost, integrated OTN switching will be more ubiquitous and touch more traffic."

THE IDEAL PLATFORM FOR TODAY: INTEGRATED OTN SWITCHING+WDM

New integrated systems that combine WDM optics and OTN electrical switching without any system density compromise are changing long-held architectural assumptions in favor of integration—first for OTN switching and WDM, and later for MPLS.

We've already discussed the benefits of decoupling client and line optics and moving away from the inflexibility and cost of muxponder architectures. But how much crossconnect switching is required in a network?

Ideally OTN switching should be integrated into every node in a core or regional transport network. This enables maximum transport efficiency by making it possible and easy to completely fill the 100Gbps, 200Gbps, or 1Tbps trunk lines and eliminate as much stranded bandwidth as possible, especially true when compared to using muxponders. Such a configuration also creates the required connection paths between nodes, giving the control plane meshiness—the least number of constraints to meet connection requirements.

Historically most hardware manufacturers divided the transport functions and switching functions into two separate hardware platforms. This creates additional cost to connect the two adjacent systems with short reach optics, significant duplication of hardware functions, and additional opex due to the use of multiple platforms along with the additional space and power these consume. These costs tended to drive solutions that took advantage of economies of scale, resulting in very large switches in a handful of nodes. This kept costs down, but limited the nodes where the benefits of switching are available, and tended to shift costs to higher layers of the network.

Putting both transport and switching into the same hardware would result in the best solution but combining the two functions historically resulted in compromise. The biggest problem has been the mismatch between short reach gray client optics density and long haul WDM optics density. Traditional discrete WDM optics use more board space and burn more power than compact XFP or SFP client optics. A chassis capable of multi-Terabit switching—if all client gray optics are used—might see effective total capacity cut 30–50% if WDM optics are deployed. Because of this carriers such as AT&T took the expected route—they bought SONET switches with 100% short reach optics to achieve maximum density and pushed the WDM transport functions externally. These same problems were evident again with the introduction of OTN switching in the core and regional networks.

An ideal platform integrates WDM transport and switching functions, connected via a more cost effective electrical backplane, and eliminates the short reach optics and duplicative hardware and opex. If DWDM transport and switching functions can be combined without compromise at the right cost, integrated OTN switching will be more ubiquitous and touch more traffic. OTN switch solutions with the highest WDM line card density—specifically those that have no density penalty versus short reach client line cards—resolve the mismatch issue that forces the equipment to be separated. Matching client and line WDM density is a critical milestone that would allow carriers to rethink the traditional hardware partitioning of transport and switching.

TOO SOON FOR PURE MPLS OR "OTN+MPLS" SWITCHING

Some in the industry have discussed migrating directly to MPLS switching as the sub-wavelength grooming/bandwidth management function, completely eliminating the OTN switching layer. Based on conversations with service providers, it seems unlikely that a pure MPLS solution would be used for this function for a number of reasons.

1. MPLS data and optical transport functions are viewed as separate hardware functions in the core network, often with separate organizations within service providers responsible for procuring and operating them.
2. There are fundamental transport OSS, processes, training, and skill sets in place that need to be updated or replaced for MPLS.
3. MPLS standards for transport such as MPLS-TP have not made as much progress as the industry hoped and will require a few more years of work.
4. Pure MPLS provides no means to partition, share, or protect a transport network that contains transparent wavelength services.

As a result, architecture approaches that have advocated the elimination of OTN switching altogether have seen little traction to date outside a small circle of specialty operators with straightforward all-packet networks with no need for multiple services.

Another approach that is being advocated is combining OTN and MPLS switching into the same core and regional WDM transport platform. This vision was amplified by Verizon and is referred to as the packet optical transport network (P-OTN). These same four issues create a similar long-term challenge in this approach, but there are also near-term issues with integrating Ethernet/MPLS and OTN in the same hardware at multi-Terabit scale.

Though new universal protocol switching chipsets coming to market may radically reduce the cost and power footprint of switching OTN and MPLS traffic in the same fabric, problems still exist on the line card. MPLS functions such as buffering and sophisticated packet lookup/forwarding decisions on the line card are still too complex to be introduced to an OTN system today without compromising the OTN switching density of the system.

Exhibit 6: Integration Without Compromise

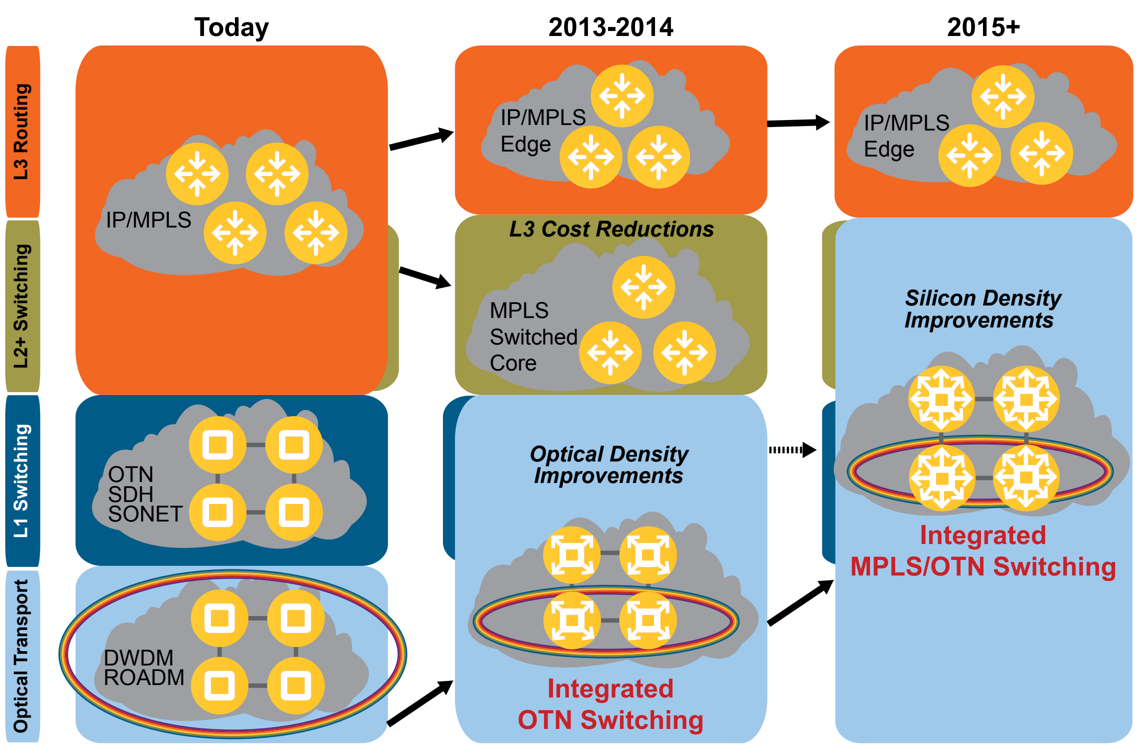

"Integration without compromise will allow OTN switching to be pervasively deployed throughout the network."

One could build a combined OTN and MPLS platform, but the density and cost would be constrained by the components and power required for MPLS functions. OTN line cards can be smaller and cheaper than an MPLS capable card, since the header processing requires buffering and specialized NPUs and memories. The resulting combined machine would lack the density of pure OTN switching platforms and cost more, resulting in carriers taking the historical route of separating these functions.

Within a few years, silicon density improvements will allow the combination of WDM, OTN switching, and MPLS label switching in large-scale platforms designed for core and regional networking. In the meantime, some vendors have adopted new hardware architectures based on switch fabrics that can eventually support both OTN and MPLS once the line card density and cost issues are conquered. We believe that transport teams that deploy these forward-looking architectures will be better prepared for deploying MPLS transport/switching alongside OTN transport/switching.

THE OPTICAL REBOOT TRIGGERS INTEGRATED OTN SWITCHING DEPLOYMENT

Carriers view the transition to 100Gbps coherent networking as a once-a-decade opportunity to reboot network architectures. Operators are using this opportunity to introduce meshed networks via integrated OTN switching, and with it, new provisioning, restoration, and utilization efficiencies.

OTN switching provides an automated way to manage the mismatch of client services that are predominantly 10Gbps and below, with the rapidly scaling WDM line side capacity that is moving to 100Gbps and beyond. Integrated OTN switching also can bring the popular concept of virtualization from the data center to optical transport networking by separating the client interfaces from a common WDM transport resource pool. In doing so, it allows more efficient use of these resources by maximizing utilization and moving beyond the architectural constraints of muxponder architectures.

Our interaction with carriers illustrates that most operators plan to deploy OTN switching, and those carriers that do account for almost all of global service provider capex. Though a few carriers plan to eschew OTN switching, we believe they have unique networks not representative of the mean.

Though the benefits of mesh based networking are clear, hardware limitations previously forced compromises and resulted in sub-optimal network architectures to reduce first year capex costs. New integrated systems that combine WDM optics and OTN switching without any compromise in terms of system density are changing long-held architectural assumptions in favor of integration—first for OTN switching and WDM, and later for MPLS.

Integration without compromise will allow OTN switching to be pervasively deployed throughout the network, dramatically increasing optical network efficiency and providing an optimal foundation for virtualized transport networks with the lowest lifetime cost of ownership. ■

WHITE PAPER AUTHOR

Andrew Schmitt
Principal Analyst, Optical
Infonetics Research
andrew@infonetics.com | +1 (408) 583.3393 | twitter: @aschmitt

ABOUT INFONETICS RESEARCH

Infonetics Research is an international market research and consulting analyst firm serving the communications industry since 1990. A leader in defining and tracking emerging and established technologies in all world regions, Infonetics helps clients plan, strategize, and compete more effectively.

Services

- Webinar, Conference, and Event Speaking
- Custom Brand and Demand-Side Market Research
- Custom Market Size and Forecasts
- Technology and White Papers
- Competitive Analysis and Due Diligence
- Market Share, Market Size, and Forecasts
- Enterprise/SMB and Service Provider Survey Research
- Continuous Research Services
- Service Provider Capex and Subscriber Analysis and Forecasts
- Consulting, Retainers, and Quick Consults

Coverage Areas

- Mobile and Wireless
- Small Cells and WiFi Offload
- Mobile Backhaul and Microwave
- Service Provider VoIP and IMS
- Broadband Access
- Video
- Enterprise Voice, Video, and UC
- Overall Market Trends and Forecasts
- Next Gen OSS and Policy
- Carrier Routing, Switching, and Ethernet
- Optical
- Service Provider Outsourcing
- Service Provider Capex and Subscribers
- Data Center and Cloud
- Security
- Enterprise Networking

SALES

- **Western North America, Asia Pacific, Latin America**: Larry Howard, Vice President, larry@infonetics.com, +1 408.583.3335
- **Eastern North America, Texas, and the Midwest**: Scott Coyne, Senior Account Director, scott@infonetics.com, +1 408.583.3395
- **Europe, the Middle East, and Africa**: George Stojsavljevic, Senior Account Director, george@infonetics.com, +44 755.488.1623
- **Japan, South Korea, China, Taiwan**: http://www.infonetics.com/contact.asp