



THE ELASTIC NETWORK

IP Over DWDM

IN A POST-SONET/
SDH WORLD



WHITEPAPER

IP OVER DWDM IN A POST-SONET/SDH WORLD

IP over DWDM, also called IP over Glass, was an exciting topic in the early 2000s as DWDM started to take off as a networking application in conjunction with the emergence of all-IP networks as an alternative to traditional SONET/SDH networks. Since then, the furor surrounding IP over DWDM has died down as the telecommunications crash caused many innovative network concepts to be abandoned. Now that the telecommunications market is seeing a new renaissance, largely due to the demands of data center operators and webscale content providers, IP over DWDM has started reappearing in the conversation, but in ways that differ significantly from the conversations at the turn of the millennium.

HISTORY

WHY? IP over DWDM was originally proposed as a way to eliminate the inefficient and costly SONET/SDH layer on networks that were carrying exclusively data (IP) traffic. This approach had merit, as SONET/SDH networks were designed to carry voice traffic and were badly scaled and over-engineered for networks based on data.

HOW? Originally, IP over DWDM proposals centered on using external transponders in DWDM equipment to convert wavelengths, as pluggable optics were not yet ubiquitous and fixed DWDM optics in routers did not (and do not) make sense economically.

The WHEREFORE? The primary benefit was eliminating the inefficient and costly SONET/SDH intermediate layer. The strongest arguments against IP over DWDM centered on the inability to route wavelengths (as ROADMs did not yet exist), forcing all data traffic between locations to be on the same route, which seemed like a significant limitation at the time.

RELEVANCE? As data traffic started to predominate in networks, SONET/SDH deployments started to decline. SONET/SDH was not designed to carry packet traffic efficiently, and alternatives such as OTN, Carrier Ethernet, and MPLS began to dominate the network transport layer.

In a world where SONET/SDH no longer exists as a layer to be avoided, what does it mean when technologists talk about IP over DWDM?



NEWS ANALYSIS, LIGHT READING 4/21/2000

Ciena Corp. <http://www.ciena.com> and Juniper Networks, Inc. <http://www.juniper.net> announced yesterday that they have successfully completed interoperability testing between 10Gbps (OC-192c/STM-64) optical and IP networking platforms. The two vendors demonstrated direct connectivity between Ciena's MultiWave CoreStream optical transport system and Juniper's M160 core router delivering 10 Gbit/s IP services over glass.

Sounds technologically fascinating. But is it practical? Many experts say no. "There are few systems if any that can put IP over WDM optics," says Scott Clavenna, principal analyst at Pioneer Consulting <http://www.pioneerconsulting.com>. "But the trend is moving away from efforts to put IP directly on glass. It's more impressive as a demonstration than anything else."

At the time, it seemed like the technology could eliminate the need for transponders, thereby reducing the cost of IP over DWDM.

But the network is still the main bottleneck. Itches from vendors like Sycamore Networks Inc. <http://www.sycamore-networks.com> to act as middlemen between IP routers and DWDM systems.

IP Over Glass,

Who Cares?



TODAY, OTHER TECHNOLOGIES HAVE GROWN IN POPULARITY



Telecommunications technology has changed quite a bit in the last decade. The advent of ROADMs means that arguments about not being able to route wavelengths are no longer valid. IP over DWDM wavelengths can easily be added, dropped, and switched around networks flexibly. Additionally, in the age of huge data center interconnect demands, the idea of forcing a great deal of data traffic to travel on the same path is not necessarily a limitation. Many data center interconnections are simple, point-to-point connections with hundreds of gigabits per second of capacity required. Finally, pluggable optics (DWDM optics or gray optics) are now the standard at speeds of 10Gbps or less, and are rapidly taking over at 100Gbps speeds.

In fact, many of the networks deployed today look exactly like the proposed, futuristic IP over DWDM networks of the past. Routers are now very often connected directly to transponders and the idea of putting a SONET/SDH network in between the router and the optical network would never be considered a valid option. In some cases, operators and vendors have even removed the DWDM equipment and put DWDM optics directly into routers and switches, a network architecture that has benefits in specific applications. IP over DWDM, rather than being an archaic concept from the history books, is exactly what is being deployed in many networks today. New optical technologies mean that there are now multiple flavors of IP over DWDM to choose from, each appropriate for a different network design.

OPTICAL TRANSPORT NETWORK (OTN)

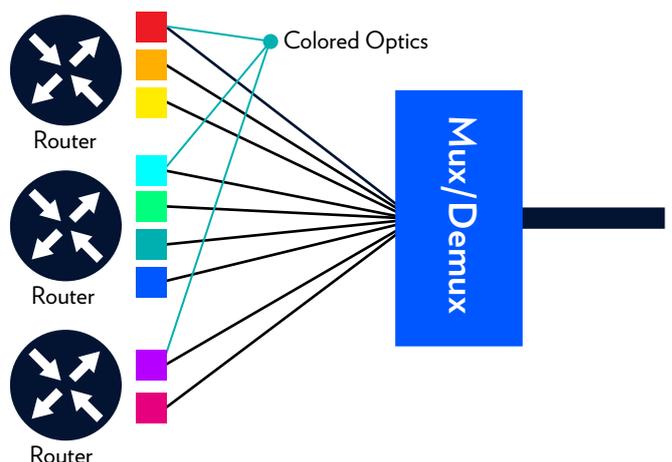


OTN has emerged as a powerful and interoperable successor to SONET/SDH that is designed better for data-centric networks. The forward error correction (FEC) built into OTN is especially useful as it allows signals to travel over longer distances without regeneration. For primarily that reason, OTN became the main Layer 1 technology deployed in metro and long-haul networks. Today, even data-only networks that aren't using the switching capabilities of OTN will use OTN-framed signals as a way to gain span budget, and it is certainly valid to claim IPoDWDM functionality in a network that deploys OTN framing. Adding OTN switching increases the flexibility and efficiency of networks, especially those carrying many different types of services.

DIRECT OPTICS



DWDM optics plugged directly into routers and switches, are, for some people, the purest definition of IP over DWDM. The router or switch has direct access to the optical layer without the need for any additional equipment other than passive multiplexers. The benefits of this architecture are reduced cost of equipment and fewer boxes to manage.



For short distances (80km or less), optics directly in routers can be a very economical option. Microsoft's 2016 announcement of a partnership with Inphi for 100Gbps pluggable optics was based on the idea that WDM optics directly in a router can save capital dollars for links of 80km or less. Optics directly in routers are a viable option for a simple Ethernet-only network, with traffic traveling between two fixed points and operated by a company with trained staff and IT/optical expertise. For networks that meet the other criteria but extend over longer distances, there may still be solutions that make sense based on pluggable optics that incorporate OTN wrapping of the data traffic so that the output from the switch can be sent over a longer distance or incorporated into an OTN-based network without the need for transponders.

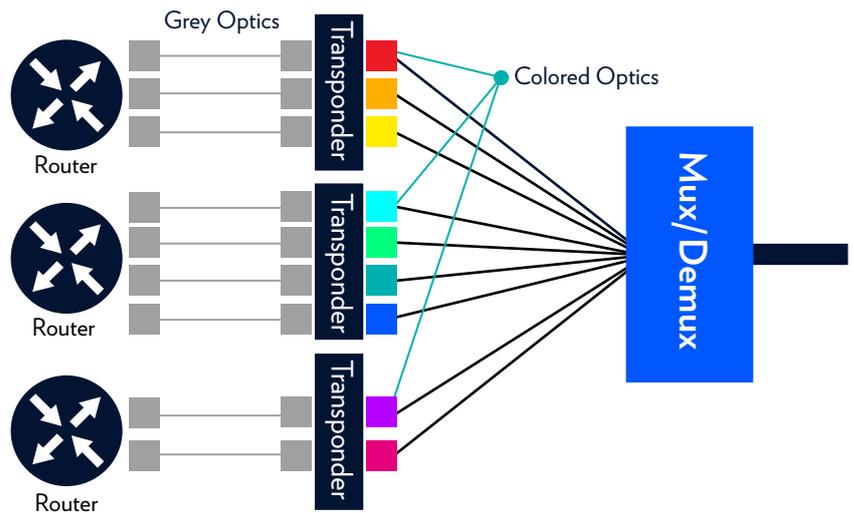
However, putting DWDM optics directly in a switch has some significant limitations, which have led most traditional carriers and many non-traditional companies (e.g., Google) to declare that they will not take this approach

NETWORK SEGMENTATION



The alternative to direct optics is network segmentation, whereby transponders are used to segment the routing part of the network from the optical/DWDM part of the network¹. Low-cost, non-DWDM “gray” optics are deployed in the routers and conversion to DWDM wavelengths is carried out using transponders on a separate shelf. The advantages of this architecture are many, including router faceplate density, fault isolation, expanded optical capabilities, and technology roadmap decoupling.

For many companies, especially those with more complex networks, a segmented network based on transponders is the right choice. Even purely packet companies like Google have publicly stated that they intend to continue building a segmented network with the transponders acting as the demarcation between the optical and routing functions. The downside of network segmentation is higher capital costs, but those costs can be offset by operational savings, router port savings, and improved network design options for more complex networks.



CHOOSING THE RIGHT SOLUTION FOR YOU

There are several tradeoffs that should be considered when determining which IPoDWDM solution is the right one for your network.



DENSITY

Faceplate density is usually the primary concern for optical interfaces in routers, especially in data centers where the majority of traffic is router-based. More interfaces on the front of the router mean more opportunities for revenue. Physically larger interfaces limit the amount of traffic that can be

¹ For simplicity and clarity, this paper does not take into account packet/optical transport platforms (POTPs) which incorporate transport at Layer 0 (DWDM), Layer 1 (OTN), and Layer 2 (Ethernet) and which may include some higher layer functionality. POTP systems are a valuable part of many networks, but are not designed to take the place of the true Layer 3 routers which are the focus of this discussion.

served by the router, regardless of the internal routing engine capacity. However, this limits the ability of routers to employ the latest generation of optical interfaces, as the latest generation optics are always larger than the last generation due to increased complexity and power restrictions.

While 10Gbps optics are now of a size comparable to 10Gbps copper, DWDM-capable 100Gbps optics have not yet reached that threshold. Instead, lower-cost and smaller gray 100Gbps optics can be used on the router to keep the faceplate density high while putting the larger DWDM-capable optics on a separate shelf. This allows the router to be used optimally without limitations imposed solely by the current generation of optics.

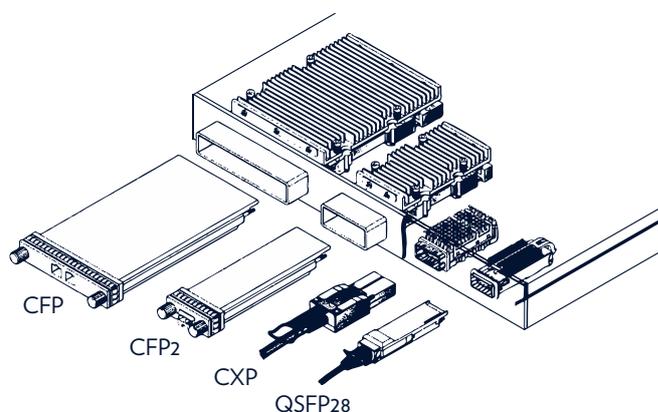
Optical muxponders offer another option for incorporating the latest optical rates with a router-based network. For example, a data center that has standardized on 10Gbps router interfaces can use a muxponder that combines multiple 10Gbps signals onto a single 100Gbps (or higher) wavelength. This eliminates the need to deploy a router or switch to combine lower speed signals and allows the use of a common, lower-cost interface throughout the data center.

The density problem will only exacerbate as optics move from 100Gbps to 400Gbps to 1Tbps and beyond. Each successive increase in speed will continue to push the limits of power and size beyond that of current optics. Keeping the DWDM optics separate from the router can help to avoid the density and capacity challenges created as optical technology progresses.



A segmented network solution would be appropriate for networks that are expected to grow over time and which need, or expect to need, optical rates higher than 10Gbps, and for which router density is a premium.

For networks with 10Gbps or lower speeds and/or short distances between sites, a direct approach might be worth considering.



The evolution of 100Gbps optics. DWDM 100G optics are moving from CFP to CFP2 while “gray” 100G optics are approaching the size of 10G SFPs.



OPTICAL CAPABILITIES

As optical technology advances and the capabilities of dynamic optical networks become critical in a modern network, routers simply do not have the built-in capabilities that are offered by dedicated optical gear. Advanced transponders not only convert the service wavelengths to DWDM, but also add the OTN framing necessary for longer distance operation and interoperability on a transport network, provide the ability to multiplex lower speed signals into higher speed signals, offer protection capabilities, and incorporate advanced optical troubleshooting tools.

While some pluggable optics offer OTN framing for longer distance operation, most transponders can do more with OTN than pluggable optics can offer. For example, transponders can provide comprehensive performance monitoring of the link at the OTN layer by looking at corrected FEC errors. This capability can provide advance information regarding the health of a connection before a failure. Transponders can also provide the ability to change the OTN FEC algorithm to optimize for distance, efficiency, or interoperability.

Modern transponders can include advanced optical troubleshooting tools such as PRBS generation and detection, OSNR measurement, OTDR capabilities, and more that make maintenance of an optical network easier

and more proactive, thereby reducing or even eliminating network downtime. By incorporating these functions directly into the transponder, the network operator can eliminate the need for external test equipment and a large staff on call to operate them. Additionally, problems in the optical layer can be identified in advance, before network failures happen, increasing overall network uptime and revenues.

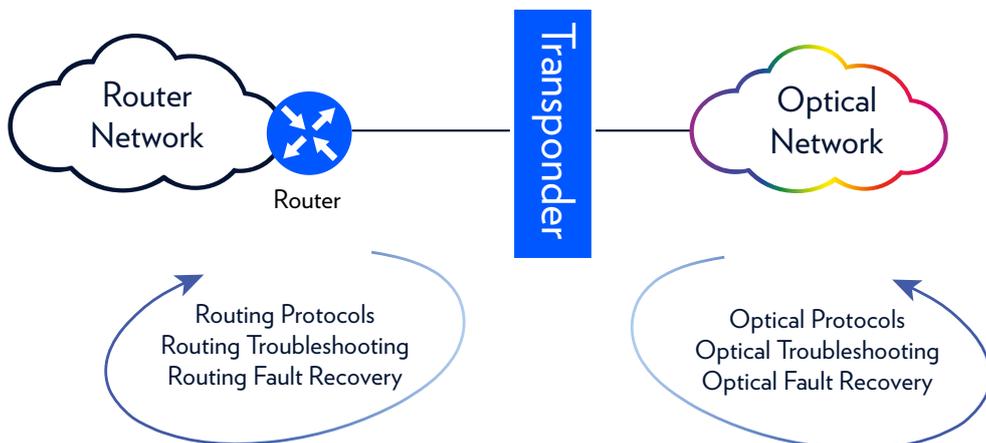


Therefore, for all but the most simple point-to-point networks, a segmented network can offer significant operational benefits that could far outweigh any cost savings from a directly integrated strategy.



OPERATIONAL DEMARCATON AND FAULT ISOLATION

Perhaps the most often cited benefit of adding a transponder layer is that transponders provide a demarcation between the switching/routing network and the optical network, which can be crucial for maintenance and network growth. Problems in the network can be isolated within the appropriate area, and upgrades to one layer of the network do not necessarily impact the other layers or other services.



In a segmented network, staff and tools that are dedicated to routing can focus on the routing part of the network while staff and tools that are dedicated to optical networking can focus on the optical part.

For larger companies with simple networks, the benefits of operational demarcation may not justify the expense of an additional transponder layer. However, most telecommunications companies have standardized on transponders based almost entirely on the benefits offered by a demarcation point between optical transport and routing.

FOR COMPANIES WITH AN IT STAFF THAT IS LESS FAMILIAR WITH OPTICAL NETWORKING, THIS SEPARATION PROVIDES A MUCH SIMPLER WAY TO TROUBLESHOOT THE NETWORK, AS THE OPTICAL FUNCTIONALITY IS ALWAYS SEPARATED FROM THE ROUTING FUNCTIONALITY. ADDITIONALLY, FAULT RECOVERY AT MULTIPLE LAYERS CAN BE MORE EASILY COORDINATED IN A SEPARATED NETWORK, AVOIDING CONFLICTS THAT CAN INCREASE NETWORK DOWNTIME.



THE SOFTWARE-DEFINED NETWORKING FUTURE



Software defined networking (SDN) is an exciting new trend in networking that is working its way into the transport part of the network. SDN offers the possibility of a truly application aware network that responds intelligently not just to bandwidth demands, but to the requirements of the applications running over the network. Multi-layer optimization, router bypass, service on demand, and other revenue-generating applications can be enabled by a SDN-based transport network.

In SDN, an intelligent orchestrator manages network configuration based on a set of standard interfaces with local controllers or equipment and in response. Currently, the plan is to have separate domain controllers for the transport layer and the routed layer, primarily because of the different sets of expertise and algorithms required to understand the technical capabilities of each layer.

The segmented layer approach is ideally designed for this type of architecture, especially if capabilities beyond simple point-to-point transport are desired at the optical layer. This is the driver behind disaggregation initiatives at some of the larger data center and content delivery network operators. However, an integrated transport capability that extends from Layer 2 packet transport down to Layer 0 WDM (e.g. with optics in a packet transport system) can also integrate very efficiently into a software-defined transport network.





CONCLUSIONS

In the new post-SONET/SDH world, IPoDWDM is still a valid network design for many data-centric networks. For some simple networks, such as intra-data center optical networks or short-haul simple interconnect, IPoDWDM may look like WDM optics plugged directly into routers. For longer distance or more complex and/or service-critical networks, IPoDWDM looks like routers connected to an external transponder network. For customers with all-IP networks, consideration of their current and planned future networking requirements will determine the optimal approach.

ECI'S IPoDWDM SOLUTIONS

Whatever solution you choose, ECI can help. ECI offers solutions that span the range of potential IPoDWDM deployments. ECI's Neptune™ system provides Layer 2 and Layer 3 capabilities with options for integrated optical capabilities that are ideal for deployment in a direct optics model. ECI's Apollo system has advanced optical capabilities that can enable a segmented network to meet its full potential, including extensive built-in test and management capabilities. The two systems can be combined with a single management address for even more comprehensive IPoDWDM solutions. Whatever your network needs, contact ECI for assistance in determining the best options for you.

Contact ECI today to learn more about ECI's IPoDWDM solutions



ABOUT ECI

ECI is a global provider of ELASTIC network solutions to CSPs, utilities as well as data center operators. Along with its long-standing, industry-proven packet-optical transport, ECI offers a variety of SDN/NFV applications, end-to-end network management, a comprehensive cyber security solution, and a range of professional services. ECI's ELASTIC solutions ensure open, future-proof, and secure communications. With ECI, customers have the luxury of choosing a network that can be tailor-made to their needs today – while being flexible enough to evolve with the changing needs of tomorrow. For more information, visit us at www.ecitele.com