

Metro WDM Network Design & Evolution: Positioning for the Transition to Optical Meshes

David W. Jenkins
Principal Engineer,
Advanced Technologies,
Tellabs

Dale A. Scholtens
Senior Principal Engineer,
Network Evolution,
Tellabs

Executive Summary

A revolution in metro transport is underway, driven by revenue and profit opportunities associated with triple-play services including voice, video and data, and demands for increasingly flexible Ethernet-based bandwidth requirements from businesses. Many service providers are considering Metro Wavelength Division Multiplexing (WDM) to meet the anticipated demand. Reconfigurable Optical Add/Drop Multiplexers (ROADMs) provide necessary flexibility for the evolution of Metro WDM networks as demand for packet-based services continues to rise. ROADMs are also one of the driving factors in the cost of Metro WDM. Service providers are therefore challenged in designing Metro WDM networks that are cost-efficient at every stage of deployment, yet nimble in light of changing service offerings and demand patterns.

An interconnected ring architecture is the most cost-efficient means of deploying and evolving Metro WDM. Interconnected rings make sense because emerging services are deep-sourced — delivered from relatively few hubbing points in a region — and therefore require more backhaul bandwidth than legacy networks. An interconnected ring architecture exploits the cost advantage of two-degree ROADM devices since it limits the need for more costly higher degree ROADMs to the role of ring interconnection. Importantly, though, it is wise to choose ROADMs with higher degrees of connectivity than immediately needed at points of current and probable interconnection, thereby leveraging the ROADM's flexibility to future-proof the metro network against unpredictable levels of demand growth.

It is more cost effective to engineer for dedicated protection of individual lightpaths in the optical layer today than for shared mesh protection. Dedicated path protection is preferred not only because of the poor economics of shared path protection in low-degree topological networks, but also because of architectural limitations with today's ROADM technology.

Even with bandwidth demand potentially exploding, interconnected ring architectures will suffice for several years. Migration from 10G (Gbps) to 40G transport should be economically feasible in the next couple of years, assuming downward cost trends on certain components continue. Service providers are therefore well-advised to deploy Metro WDM equipment that can support 40G transport, without regeneration, across the same distances that 10G can be carried today. By deploying 40G-capable equipment, service providers position themselves for increased capacity on high-demand lightpaths without forklift upgrades. Independent of 40G deployment, this paper describes a non-intuitive technique of ring

division that essentially doubles capacity in a congested region of a metro network. Having engineered extra ROADM connectivity at network junctions to begin with, service providers can use the technique to advance Metro WDM networks toward a mesh architecture. The technique adds significant capacity in congested regions of a network and requires small capital outlay for the benefit accrued. By exercising the technique several times as different areas of a network approach congestion, a service provider can gracefully and cost-effectively migrate from rings toward a mesh.

Notwithstanding the transition to 40G, service providers should therefore plan for today's Metro WDM deployments to evolve to optical meshes. By seeding with ROADMs that support high-degree connectivity as networks are initially built, the ring division technique can be applied to transform interconnected rings to meshes as demand warrants. Complemented by further advances in optical component architecture, shared protection in Metro WDM networks may eventually become cost-effective — even achieving the considerable savings in protection bandwidth observed in some SONET/SDH meshes today.

Introduction

A revolution in metro transport networking is underway, driven by revenue and profit opportunities from residential service bundles that include Internet Protocol Television (IPTV), along with increasing demand from businesses for IP- and Ethernet-based services. Many service providers are investing in multiservice access and transport networks to provide for these services, deploying equipment that supports optical interfaces exclusively.

As SONET and SDH were standardized in the 1990s, concern arose over the inefficiency of the dedicated protection schemes used in SONET/SDH rings. Indeed, the driver for the Bidirectional Line Switched Ring/Multiplex Section-Shared Protection Ring (BLSR/MS-SPRing) was a desire to improve the protection efficiency of rings as compared to the 1+1 dedicated protection defined by SONET's Unidirectional Path Switched Ring (UPSR), which allocated 50% of available bandwidth for protection. BLSR — especially its transoceanic variant — improved upon UPSR by allocating protection bandwidth just sufficient enough to accommodate traffic actually requiring protection, and making use of protection bandwidth only during fault conditions. In adopting this approach, the normally unused protection bandwidth became available for a “preemptible” class of traffic, improving the overall cost efficiency of transmission. Although BLSR's preemptible traffic class was not widely embraced by service providers for operational and marketing reasons, BLSR also

formalized the notion of Nonpreemptible Unprotected Traffic (NUT) — bandwidth that does not require protection. In practice, NUT enabled practical improvement in ring efficiency since bandwidth that would have been reserved for protection became available to carry yet more NUT traffic.

As SONET/SDH became widely deployed, rings were interconnected. It was soon realized that if networks were constructed with greater connectivity than afforded by interconnecting rings at one or two points, relatively more protection bandwidth could be shared within a network as a whole. Subsequently, there has been much research in mesh network architectures and strategies for shared protection in meshes. Indeed, some service providers made strategic decisions to build meshes to capture the savings in protection bandwidth in the last few years, recognizing that the cost of introducing the necessary dynamic switching intelligence within network elements was more than offset by protection bandwidth savings and simplified provisioning operations systems.

In these same few years, Metro WDM systems have moved from point-to-point applications for fiber relief to generalized application in Metro WDM networks. Indeed, Metro WDM is critical to the business cases of many service providers who recognize the potential explosion in bandwidth demand over the next few years. Metro WDM has become attractive due to increased levels of component integration, exemplified by ROADMs that incorporate variable optical attenuation, and electrical fabrics capable of grooming and switching both packet and TDM flows at rates of hundreds of Gigabytes per second. The Tellabs® 7100 Optical Transport System (OTS), for example, can carry 44 10 Gbps wavelengths for nearly a thousand miles without 3R regeneration, and is architected to support 40G transmission in the future. By exploiting the system's sub-wavelength switching, service providers can often justify deployment when as little as 30 or 40 Gbps of capacity is needed in an area, knowing that the system's ROADM-based optical switches provide flexibility for ongoing network expansion.

We discuss an approach to Metro WDM design that exploits the ROADM's flexibility in order to minimize capital outlay for WDM infrastructure over time. In addition, the availability of multi-degree ROADMs resurrects questions around the efficacy of optical shared mesh protection; we therefore discuss the feasibility of various protection schemes in light of ROADM architecture and overall network economics.

The ROADM Revolution

A two-degree ROADM subsystem is shown in Figure 1. The subsystem comprises two ROADM devices connected back to back. It is termed a two-degree subsystem by virtue of supporting two WDM line interfaces (East and West). Just as with electrically-based add/drop devices, a series of two-degree ROADMs can be joined to form a ring.

While the architecture of different vendors' devices varies, the following connectivity is available:

- Adds and drops are done in association with the WDM interface with which a ROADM affiliates. For example, a transponder connected to one of the add/drop port pairs of the West ROADM device in Figure 1 can launch its signal toward the West interface, and can terminate a signal from the West interface. It cannot, however, launch or terminate in association with the East interface.

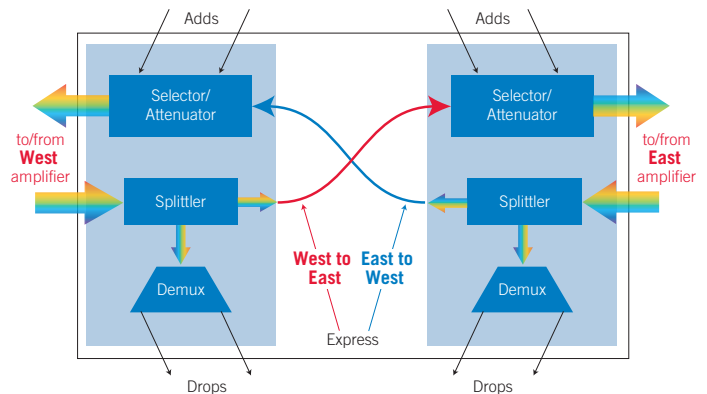


Figure 1. Two-degree ROADM subsystem

- Pass-through is effected via the splitter and selector/attenuator functions. The optical power of each wavelength entering from the West amplifier is divided by the West ROADM device's splitter. Some power is diverted to the West drop port in case the signal is being dropped, but the remaining power is forwarded to the Selector/Attenuator in the East ROADM via the Express Interface linking the West and East ROADM devices (red line in Figure 1). The East ROADM device's selector/attenuator can be dynamically configured to forward wavelengths received from the express interface toward the East interface, or instead insert wavelengths from its own add ports. The attenuation function assures that all wavelengths are at appropriate power levels in relation to one another before being amplified by the East amplifier.
- ROADM devices that position a splitter function ahead of the selector/attenuator function (which is the case shown in Figure 1) also support drop-and-continue connections. A wavelength entering from the West amplifier can be simultaneously dropped by the West ROADM and forwarded through the East ROADM device to the East interface. Drop-and-continue connectivity is cost-effective for point-to-multipoint signals such as those associated with broadcast and pay-per-view IPTV; it minimizes transport bandwidth and reduces the amount of higher level switching needed for these services.

The Cost of Connectivity

Many service providers commencing WDM rollout in a region begin simply, with a single ring traversing the set of sites from which services are initially to be provided. However, Metro WDM networks can grow to connect cities and towns that may be hundreds of miles apart. Therefore, capability to extend the initial ring in multiple directions is needed. ROADMs provide this capability.

ROADMs supporting up to eight-degree connectivity are available today, and ROADMs supporting ten degrees are on the horizon. These devices enable rings to be interconnected, and of course the ability to construct WDM meshes. Higher-degree ROADMs also give service providers the freedom to rearrange network topology over time as different services are rolled out, and to react swiftly to changes in service demand patterns.

However, higher-degree devices are limited by internal losses attributable to the splitter and selector functions, as well as the increased complexity associated with selector/attenuator function. This is evident even with a four-degree ROADM subsystem, illustrated in Figure 2. Because of the number of express interfaces increases geometrically with the degree of connectivity, ROADM cost also increases with subsystem degree.

Why does this matter? The cost of a WDM network is dictated by its major media plane components. Broadly, these are:

- *WDM spans (amplifiers)*: WDM is attractive because its amplification costs are lower than the corresponding costs of stacked SONET/SDH multiplexers to carry equivalent bandwidth. However, the decision to deploy a WDM link must be taken carefully because the stepwise cost is still significant. Service providers should therefore leverage existing capacity in their WDM networks as effectively as possible before deciding to add links.
- *Optical switching (ROADMs)*: As suggested above, the flexibility afforded by ROADMs comes at a price. To control network costs, then, service providers must strike a balance. Higher-degree connectivity is clearly justified at major hubs and interconnection points, and low degree connectivity is generally all that is necessary in outlying areas. But in a metro network there is often a foggy area for which the appropriate degree of connectivity is debatable in light of forecasting uncertainty involving service mixes, service take-up rates and the like.
- *Optical adaptation (WDM transponder lasers)*: Transponder costs closely follow the addition of traffic to the network, and are a lesser consideration since their costs more directly align with service orders. Still, design and evolutionary strategies that minimize the number of transponder lasers — for example, by reducing need for regeneration — tend to minimize the cost of a network overall.

Tellabs has been engaged many times to plan Metro WDM network rollouts that are capital-efficient, achieving the balance that service providers seek. We consistently observe that the lowest cost approach to Metro WDM is to construct rings traversing many service hubs and Central Offices (COs) in a region, with local distribution from these rings as necessary. This is perhaps counter-intuitive given the savings attributed to SONET/SDH meshes over the last several years. However, our modeling shows that higher-degree ROADMs allow for better utilization of WDM links in architectures of interconnected rings (network degree closer to 2.0 than, say, 3.0) as opposed to mesh designs (say, degree approaching 3.0 or higher) by deferring the need for additional amplifiers and even entire WDM systems in time.

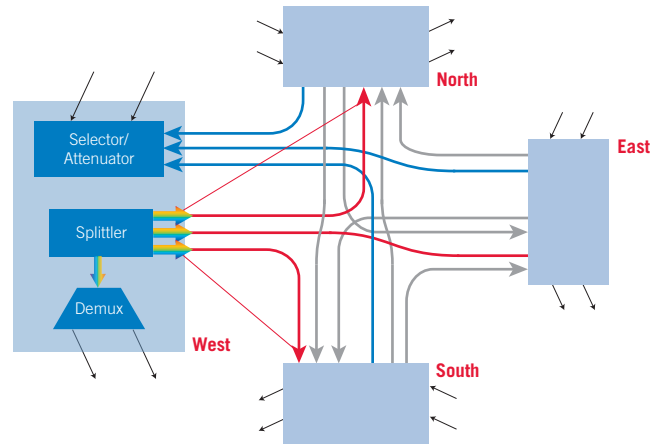


Figure 2. Four-degree ROADM

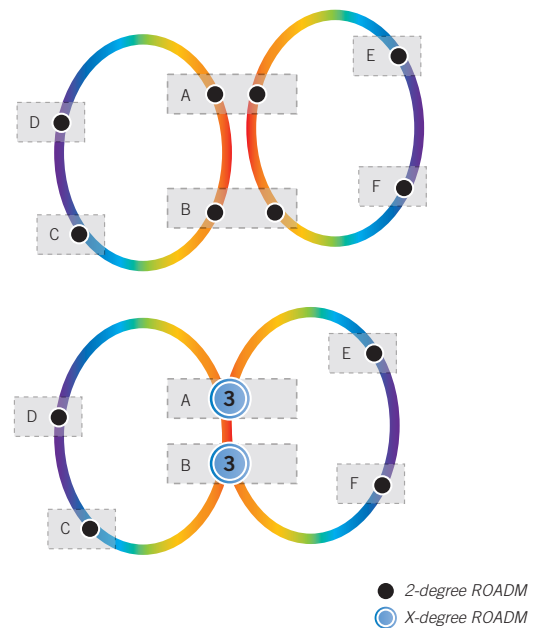


Figure 3. Using ROADMs to defer WDM links

This approach can be appreciated by comparing the costs of two independent WDM rings that traverse a pair of locations in common, versus the same pair of rings constructed using a common WDM link between those locations, as illustrated in Figure 3. In this case, the cost of the three-degree ROADMs trades favorably against the cost of a pair of two-degree ROADMs and an additional link between sites A and B. First, the three-degree ROADMs eliminate cost associated with interconnecting optical demands between the rings within those sites. Second, and more importantly, the higher-degree ROADMs defer the expenditure for the pair of amplifiers on the second WDM link between sites A & B until the capacity of the

shared link is exhausted. Given typical mixes of protected and unprotected demands, this does not occur until both rings are at least half full.

In modeling large networks, we have seen that mesh architectures exhibit the same deficiency as independently engineered WDM rings. When lightly loaded, meshes require relatively more links to service a given demand set than interconnected rings, and WDM links have associated costs. Therefore, the time-weighted cost of a mesh — especially at early stages of Metro WDM network build-out when few demands need to be accommodated — is higher than the cost of a network of rings interconnected by multi-degree ROADMs. The disadvantage of the mesh is not overcome until links become heavily used, which occurs only over time. The situation can be appreciated by considering a group of, say, 8 nodes requiring protected services. When the number of lightpaths needed is few, the minimum cost to link the nodes is a single ring comprising 8 links. Any other structure will require more links, hence greater capital outlay early on.

Clearly, a ROADM-based deployment strategy involving interconnected rings makes sense, but exactly how much connectivity should a service provider allow for over time? For example, Figure 3 depicts three-degree ROADMs. But four-degree ROADMs make more sense if it is anticipated that the link between sites A and B will in fact eventually block, since the second WDM link can be added between sites A and B without impacting existing traffic.

ROADM vendors are of course focused on reducing costs of their devices. However, relatively more attention is being given to two-degree devices than devices of higher degree. We foresee that ROADM subsystems will track along a knee-shaped curve suggested by Figure 4. This cost trend reinforces the case for interconnected rings, and in fact argues that service providers should opt for high-degree connectivity at ring interconnect points (and other probable points of network expansion) as a hedge against uncertain demand forecasts.




Figure 4. ROADM cost profile

Such a strategy nicely balances ongoing capital outlay with long-term flexibility. Though it requires some number of higher-degree ROADMs during build-out, it defers most expenses for WDM links until they are needed and enables network expansion without service interruption.

Implications of Deep-sourced Services

Metro WDM bandwidth demand is primarily driven by demand for new services, for example, as IPTV rollout takes place throughout a region, or as businesses increasingly demand Ethernet and IP-based services. Service providers intent on shuttering COs may also strategically migrate portions of existing transport networks onto their WDM infrastructure to minimize operational expense. As the metro network expands, many service delivery points become increasingly deep-sourced, and there is corresponding need to aggregate traffic toward relatively few service delivery points. The effect of deep sourcing is apparent, for example, with the transition from TDM voice to VoIP: The number of switching points in a region becomes far fewer, but the amount of traffic aggregated toward each point increases.

But voice traffic pales in comparison with the bandwidth needed for other services. On the consumer front, as IPTV transitions from a broadcast to a unicast (video on demand) model the bandwidth associated with the service may grow by two orders of magnitude, from a few Gbps per Video Serving Office (VSO) to well over a hundred Gbps per 10,000 subscriber VSO.¹ Even basic Internet service is on a speed ramp driven by competitive forces and visually rich, Web-based applications: Internet edge routers handled traffic predominantly sourced from 56 kbps access circuits just a few years ago, but 3-6 Mbps access is now common, and 30-100 Mbps service is appearing as fiber is deployed deeper and deeper in access networks.

On the commercial front, businesses increasingly seek higher-bandwidth WAN connectivity, not only to facilitate day-to-day communications between sites but also, for example, to maintain mission-critical storage. In the United States, law and policy concerning record retention and disaster recovery — such as Sarbanes-Oxley, HIPAA, and SEC 17a-4 — are furthering the demand. Enterprise Strategy Group estimates that digital archive capacity will grow 90% year-on-year worldwide through 2010 to total 28,000 Petabytes.² This rate of accumulation is stimulating demand for managed storage and archival solutions, creating significant opportunities for Storage Service Providers (SSPs). SSPs tend to deploy just a few server sites in metropolitan regions, in turn creating a need for aggregation and backhaul to those sites.

¹ Assume each subscriber has one High Definition TV (HDTV) and two Standard Definition TVs (SDTV). Assuming advanced coding, HDTV requires about 9Mbps and SDTV requires 3Mbps. In a pure broadcast model for which 150 SDTV and 40 HDTV channels are delivered, $(150 \times 3) + (40 \times 9) = 810\text{Mbps}$ is delivered independent of the number of subscribers. In a unicast/VoD model, we have $(10,000 \times (9\text{Mbps} + (2 \times 3\text{Mbps}))) = 150\text{Gbps}$.

² See: Web summary of report entitled "Digital Archiving: End-User Survey and Market Forecast 2006-2010," Enterprise Strategy Group, January 2006.

Even traditional private line services are arguably deep-sourced (or, at least “deep-connected”) as are the SONET/SDH network-builder services used internally to service provider networks. Emerging Ethernet private line and protocol-independent wavelength services are also deep-sourced. In practice, the routing of such services tends to be through preferred hubbing points as a matter of policy — for example, through sites with test heads or higher order multiplexing facilities.

The implication of deep-sourcing is straightforward: Once an architecture is decided upon, it is unlikely that the degree of connectivity in a given area of the metro network needs to change significantly over time since new services are likely to be sourced through the same hubs as existing services. Even if a new hub is added, it is likely that the hub will be positioned along an existing ring. Thus, a Metro WDM network initially laid out as interconnected rings is to likely propagate along the rings since this requires only the addition of lasers as bandwidth demands increase. Still, higher-degree ROADMs offer service providers the flexibility to migrate toward meshes if required, and so are a cost-effective hedge for network evolution.




Figure 5. Metro WDM network of interconnected rings

ROADM Support for Optical Layer Protection

Figure 5 depicts an Metro WDM network designed for deep-sourced services, comprised of interconnecting rings (degree approximately 2.4).³ While the 1+1 path protection schemes of SONET/SDH immediately come to mind, it is useful to evaluate whether shared optical protection can be more effective than dedicated schemes, along with the feasibility of implementing such protection within ROADMs. To assess, consider the following:

- *Shared protection in individual rings.* Shared protection is effected by either turning traffic around at a point of failure (like BLSR/MS-SPRing), or redirecting it from its source upon failure (transoceanic BLSR). As previously described in association with Figure 1, a wavelength arriving on a given WDM interface can either be dropped or expressed along to another WDM interface, but it cannot be returned to the interface on which it arrived. Therefore, BLSR-like optical protection is infeasible with today’s ROADM architectures. However, even if ROADMs could cost-effectively be re-architected to support turnaround, the network-level benefit would be minimal: The efficacy of turnaround in low-degree networks is poor, nominally around 5-6% savings in distance-bandwidth product.⁴
- *Shared protection network-wide.* Referring again to Figure 1 and its associated text, wavelengths added/dropped by a given ROADM device affiliate with that device’s WDM interface. In the general case of an N-degree ROADM subsystem, a given laser/photodiode can launch/terminate a signal on only one of the subsystem’s N WDM interfaces. Consequently, the ability to provide shared protection is compromised. It can be supported along the interior of a lightpath — where light both enters and leaves a node on WDM links, and




Figure 6. SNCP-protected lightpath

therefore traverses the express interfaces. But shared protection is not possible on the first and last links of a lightpath since the added/dropped wavelength does not traverse an express interface. To fully support shared protection requires that added/dropped wavelengths be de-coupled from individual ROADM devices — which requires an additional optical matrix between transponders and the ROADM device array, or its functional equivalent. Thus, even though shared protection has been shown viable in networks of suitable degree, the cost and complexity of further ROADM architectural changes must be traded against the worth of a fully functional shared protection architecture.

³ The degree of a network is the average of the degree of the network’s nodes.

⁴ See: Wayne D. Grover and John Doucette, “Design of a Meta-Mesh of Chain Subnetworks: Enhancing the Attractiveness of Mesh-Restorable WDM Networking on Low Connectivity Graphs,” IEEE Journal on Selected Areas in Communications, Vol. 20, No. 1, January 2002, pp.47-61.

- Perhaps more significantly, past analyses ascribing value to shared mesh protection in SONET/SDH networks presume that each node is capable of both space and time switching. In the WDM realm, ROADMs provide space switching alone. The second degree of freedom in WDM is frequency switching — changing the color of light along its path. Today, there is no way to do frequency switching cost effectively other than with back-to-back transponders. However, their connectivity is limited to designated pairs of WDM interfaces because of the ROADM device limitations described previously. Significant advances in optical technology are needed to provide the needed two-dimensional flexibility — arguably a blending of semiconductor laser and ROADM technologies. Again, such changes will come only with additional component complexity (hence, cost) and must be weighed against gains attributable to shared protection.
- *Dedicated path protection.* UPSR and SDH's Sub-Network Connection (SNC) protection are reference techniques for path protection. UPSR and SNC both require 1+1 signal bi-cast at the transmitting edges of the protection domain, along with corresponding selector functions at the receiving edges. UPSR protects a path within individual rings, while SNC protects the subnetwork as a whole. Techniques analogous to UPSR and SNC can be used for individual wavelengths in the WDM domain.⁵
- ROADM devices inherently support the bi-cast functionality necessary for UPSR-style optical channel protection. However, experience with dual-homed SONET rings has revealed the complexity of selection across such interconnects, and the difficulty of managing bandwidth along the common link between sites (for example, the link between sites A and B in the lower half of Figure 3) is well-known.
- *In WDM networks, SNC is preferred.* A demand protected by SNC is illustrated in Figure 6. SNC effectively creates a two node ring for each protected demand. In contrast to UPSR, SNC's protection structure is independent of the underlying WDM link topology, meaning that the underlying WDM infrastructure can evolve independent of the provisioning of its optical clients. SNC requires only that diverse routing be possible through the network, and that a single bi-cast/select function be provided at each endpoint (at A and B in Figure 6). There is no difficulty crossing ring interconnects, as is the case with UPSR.
- SNC is also amenable to automation using standard routing protocols such as OSPF. Routing protocol-based techniques can handle not only demands protected within the WDM network, as illustrated, but also client-protected demands (for example, at the IP layer), thereby affording protection for lasers at either end of a subnetwork connection without increasing the complexity of provisioning operations.

Before leaving the topic of protection, it is useful to consider the strategy for interconnecting WDM rings. Importantly, dual interconnection comes at no cost. In Figure 6, bandwidth demand to the core ring through point C cannot be protected on a node-diverse basis. For the topology illustrated, the particular ring involved could have been arranged to connect the core ring at both nodes C and D, not just C alone. The costs are the same; six ROADM devices and associated amplifiers are used whether the rings are singly or doubly connected (i.e., $4+2 = 3+3$).⁶

Leveraging Interconnected Rings — The Stealthy Path to Mesh

Metro bandwidth demand will inevitably increase. Because services are deep-sourced, it is likely that congestion will first occur in the core of the network — perhaps on a ring connecting major cities and towns in a region. Congestion manifests as wavelength blocking — where wavelength assignments for existing lightpaths are such that a new traffic demand cannot be carried through the network on light of a single color.

Service providers have several choices when encountering congestion, with relief available at both lightpath level and WDM link level as indicated in Table 1. In choosing amongst relief strategies, service providers desire to minimize capital outlay; at the same time they want to be confident that expenditures are appropriate in context of longer term network evolution.

Technique	Network CapEx Efficiency
<i>Lightpath Based</i>	
Reroute or change color of existing lightpath	Ideal: Zero cost (but seldom possible)
Insert regenerator	Poor: One demand served at cost of two transponders
Upgrade speed	Better: 4x capacity for 2.5x cost (in near future)
<i>WDM Link Based</i>	
Bypass regional ring	Poor: Viable only if large number of demands arise between two edges
Bisect regional ring (add one WDM "shortcut")	Poor: Provides for small to moderate number of additional demands
Divide regional ring into two interconnected rings (add two WDM links)	Excellent: Near doubling of core network capacity

Table 1. Efficiency of various congestion relief approaches

⁵ Refer to ITU G.872 for a formal definition of SNC. For purposes of this discussion, we restrict ourselves to 1+1 SNC schemes.

⁶ To be fair, it may be impractical to provide dual interconnect everywhere for reasons of geography or lack of sufficient demands requiring true node-diverse routing. In such cases, there is a small compromise in overall service availability to traverse a single interconnect node.




Figure 7. Ring bypass and bisection




Figure 8. Regional ring division

At the lightpath level, costs of congestion relief can be fairly well synchronized with increases in demand. Options at this level are:

- *Rerouting and/or changing color.* In a few cases, relief can be obtained at no cost by changing the routing or color of an existing lightpath in order to free resources for a new demand. Importantly, the ROADM's flexibility coupled with control-plane based approaches to lightpath provisioning⁷ makes reroute and color change more feasible than has historically been the case. Using the control plane, such changes can be performed quickly, and without repeated interaction with OSSs since only the final result needs to be communicated northbound. Still, this technique is very situational since an appropriate pattern of demands must already exist in the network. Service Level Agreements (SLAs) must also allow momentary interruption of service if unprotected lightpaths are manipulated.
- *Inserting regenerators.* Regeneration comes at a cost that is small compared to WDM link addition; it also allows the routing of lightpaths in different areas of the network to be decoupled. If the incidence of blocking is low, regeneration makes sense. However, overall cost-effectiveness of the network is compromised if very much regeneration is needed, since equipment slots are tied up for regeneration instead of servicing additional demand.
- *Introducing higher bandwidth lightpaths.* Moving from 10G to 40G transmission provides the most leverage, and service providers will surely exploit this approach as it becomes economically feasible. Electronically controlled dispersion compensators that allow 40G lightpaths to travel as far as 10G lightpaths are appearing today. Though not cost-effective yet, 40G devices appear to be on path for mass introduction. Coupled with cost reductions in 40G tunable lasers, we expect the historical 2.5x cost/bandwidth ratio for 40G transponders to be achieved within the next couple of years (that is, 40G price \approx 2.5 x 10G price). Pragmatic service providers will therefore avoid wavelength blocking by upgrading congested areas of their networks to 40G, freeing 10G lightpaths in the process. If

traffic is growing at 40% per year, for example, a complete 10G to 40G upgrade defers additional WDM links for nearly four years.⁸

At the WDM link level, costs of creating capacity are higher, and thus greater benefit to the network as a whole should be expected when links are added. It is unwise to add links that cannot be well-utilized very soon, or do not create significant additional lightpath capacity in the network as a whole. At the WDM link level, we consider the following options:

- *Bypassing existing rings.* Refer to Figure 7, which increases the degree of nodes A and B on the edges of the network previously illustrated. By increasing both of these nodes to degree three, a WDM link (red dashed line) can be added between them. In context of the network as a whole, bypass links such as this have little topological value. They do not prove cost-effective unless there is significant unprotected traffic between the sites they are connecting — usually many wavelengths. Otherwise, protected traffic between A and B must use the existing network anyway.⁹ Also, demands originating from nodes A and B must detour into the distribution rings in order to make use of the link.
- *Bisecting a ring.* The intuitive approach to blocking in the core of the network is to increase the degree of connectivity between a pair of sites to relieve blockage. The link between nodes C and D (blue dashed line) in Figure 7 illustrates this case. The technique appears attractive from the standpoint of cost since only a single link is being added. And it seems that by adding such links over time, the network can evolve toward a mesh. However, the network's degree of connectivity — and therefore the flexibility to route arbitrary demands — is not significantly influenced by adding just a few such links. In a network providing deep-sourced services, the practical effect of such additions is to provide about half of the anticipated bandwidth. Once again, the fallacy is apparent when considering the proportion of traffic requiring protection, since protection traffic must traverse the existing network.

⁷ A companion paper covers ASON/GMPLS control-plane based operation of WDM and sub-wavelength networks.

⁸ Whether using 10G or 40G transmission, aggregating and switching traffic to fill lightpaths completely at network junction points is critical to metro WDM efficiency. Equipment such as the Tellabs® 7100 OTS supports integrated packet and TDM switching and aggregation, typically offering 30% savings at ring interconnects compared to use of external equipment for the purpose.

⁹ It makes no difference whether the WDM network (a) provides protection and diverse routing on behalf of simplex client demands, or (b) carries pairs of client-protected (or load-shared) demands on diverse paths. Either way, pairs of information flows traverse the network diversely between pairs of nodes. Hence, freedom to create diversely routed pairs of lightpaths is a measure of a topology's effectiveness.

- *Dividing a ring into two rings.* Counter-intuitively, adding *pairs* of links between sites on a ring is actually much more cost-effective than bisection. This is shown in Figure 8, where nodes A and H have been increased by two degrees each. In Figure 8, the addition of the two links divides the core ring into overlapping upper and lower rings. The upper ring comprises nodes A-B-C-D-E-F-G-H, while the lower comprises A-H-I-J-K-L.
- By dividing a ring approaching congestion into a pair of rings, capacity becomes available on each of the newly created rings. In the example shown, the second link addition is well worth its cost, since the bandwidth of the entire core of the network is essentially doubled: If the ring being divided has the capacity of n wavelengths, the addition of the two links shown provides for n wavelengths in each of the two rings created. This doubling of bandwidth also reduces the likelihood of wavelength blocking as demands enter the core of the network; thus the cost of the second link is further offset through the reduced likelihood of regeneration for purposes of changing colors.

To summarize, lightpath-level relief via rearrangement or regeneration is appropriate when little additional demand is anticipated for some time. Where demand is increasing in a pattern consistent with existing lightpath routes, 10G to 40G upgrade should be considered to forestall relatively greater expenditure on WDM links. Service providers are therefore wise to choose WDM equipment that will support 40G upgrade without service interruption, and transport 40G lightpaths at the same distances as 10G lightpaths without having to upgrade existing ROADMs or amplifiers.

When links must be added, service providers should opt for ring division — adding pairs of links — over ring bypass and ring bisection. Apart from doubling capacity in the area of the network where the division occurs, ring division leads to a more fully connected network over time. Using ring division, service providers can be assured of providing bandwidth in areas of the network where it is most needed. Ring division also allows a graceful migration to meshed architectures. For instance, five division operations in the network of Figure 8 will change its degree to about 2.8, a level at which shared mesh protection begins to show some benefits.

Mesh — and Eventually — Shared Mesh Protection

We have shown that building a Metro WDM mesh outright is an inefficient use of capital. Service providers will therefore prefer to engineer interconnected WDM rings for some time. However, neither comfort with ring architectures grounded in SONET/SDH experience nor the transition to 40G should dissuade service providers from architecting ring-based networks for migration to meshes. In deploying Metro WDM, service providers should think about the evolved topology of the Metro WDM network, which will likely be a mesh comprising some number of high-degree nodes interconnected by chains of two-degree nodes. The key to migrating is to install higher-degree ROADMs than are immediately needed at probable points of ring interconnection — in essence, scattering space switching around at important junctions in the network, independent of whether the switching capability will immediately be used. By strategically deploying higher-degree ROADMs today and using the ring division technique we describe, service providers can gracefully evolve their Metro WDM networks toward meshes — in concert with whatever demands arise, when they arise, and with capital outlay aligned with those demands.

As a matter of policy, meshes may come sooner than suggested by bandwidth demand alone. Survivability trumps capital constraints that ordinarily guide network architecture and deployment; the recent experience with Hurricane Katrina is a reminder. Service providers may therefore choose to engineer “extra” WDM links early on to provide a slightly higher degree of connectivity than available with interconnected rings. This increases the likelihood that paths will be available for critical services, even in the face of multiple failures. The Automatic Switched Optical Network/Generalized MultiProtocol Switching (ASON/GMPLS) control plane is an enabler in this strategy. For example, a critical SNC-protected service can be provisioned by ordinary means. But in case both working and protection paths fail, the control plane can be used to quickly recalculate a temporary working route using available capacity, even preempting NUT traffic on the “extra” links if necessary.

Whether driven by demand or concerns over survivability, we can expect that Metro WDM networks initially engineered as rings will evolve to meshes. As connectivity increases, service providers will surely turn their attention to shared mesh protection. The attention is justified in the face of decreasing revenue-per-bit from packet-based services. SNC requires 100% additional bandwidth on raw basis, and sometimes over 200% in terms of distance-bandwidth product (when the comparative lengths of working and protection paths are considered). In networks with sufficient connectivity, shared protection can save 10-30% of raw protection bandwidth, and often bring distance-bandwidth product for protection down to the 60% range, or even lower.¹⁰

¹⁰ See John Doucette, and Wayne Grover, “Comparison of Mesh Protection and Restoration Schemes and the Dependency on Graph Connectivity,” 3rd International Workshop on the Design of Reliable Communication Networks (DRCN 2001), Budapest, Hungary, October 2001.



To achieve such savings in distance-bandwidth product requires components that are not yet available — components that offer the freedom to change color (and ideally, perform full regeneration) as needed along shared paths. Architecturally, this functionality must be closely coupled with higher-degree ROADM devices — even directly integrated. Today, color changers are being researched but none is near commercialization, much less integration with a ROADM. Therefore, it will be some time before such components are available.

Fortunately, service providers have time and current technology on their side. ROADMs provide the flexibility needed for present Metro WDM buildouts. Ring division using high-degree ROADMs will provide capacity for years to come, independent of the timing of the migration to 40G. As a leading vendor of optical transport gear, Tellabs will continue to influence development of WDM components — ROADMs, amplifiers and lasers alike. Service providers can therefore be assured that their Metro WDM networks will remain capital efficient as WDM technology continues to advance.

North America

Tellabs
One Tellabs Center
1415 West Diehl Road
Naperville, IL 60563
U.S.A.
+1 630 798 8800
Fax: +1 630 798 2000

Asia Pacific

Tellabs
3 Anson Road
#14-01 Springleaf Tower
Singapore 079909
Republic of Singapore
+65 6215 6411
Fax: +65 6215 6422

Europe, Middle East & Africa

Tellabs
Abbey Place
24-28 Easton Street
High Wycombe, Bucks
HP11 1NT
United Kingdom
+44 871 574 7000
Fax: +44 871 574 7151

Latin America & Caribbean

Tellabs
Rua James Joule No. 92
EDIFÍCIO PLAZA I
São Paulo – SP
04576-080
Brasil
+55 11 3572 6200
Fax: +55 11 3572 6225

The following trademarks and service marks are owned by Tellabs Operations, Inc., or its affiliates in the United States and/or in other countries: TELLABS®, TELLABS and T symbol®, and T symbol®. Statements herein may contain projections or other forward-looking statements regarding future events, products, features, technology and resulting commercial or technological benefits and advantages. These statements are for discussion purposes only, are subject to change and are not to be construed as instructions, product specifications, guarantees or warranties. Actual results may differ materially. The information contained herein is not a commitment, promise or legal obligation to deliver any material, code, feature or functionality. It is intended to outline Tellabs' general product direction. The development, release and timing of any material, code, feature or functionality described herein remains at Tellabs' sole discretion.