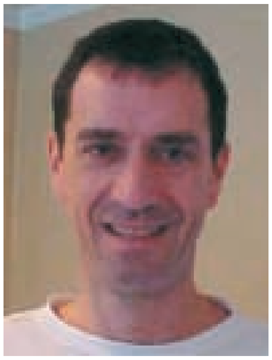

IP over Optical:

From Packets to Photons

Chris Metz • *Cisco Systems* • *chmetz@cisco.com*

The computing industry is currently experiencing a love affair with all things optical. Not since ATM's heyday in the mid-1990s has a network technology generated so much excitement, innovation, or debate—not to mention hype and venture capital. Optical networking is the subject of countless research papers; of civilized (and uncivilized) discourse in many standards bodies; and of many industry conferences. Moreover, it is the core strategic component of many traditional and start-up network enterprises. And why not? This technology promises innovative developments that will simultaneously make communications faster and less expensive.

Optical technology can also be viewed as just a means for creating high-bandwidth point-to-point connections, and a single data-link technology—no matter how fast or cheap—does not a network make. Networks are built by interconnecting collections of hosts and data links of varying sizes, shapes, and speeds to IP routers. Connecting the collections to other collections of IP routers and data links forms the collective internetwork; the Internet is the prime example.

While optical technology indeed delivers orders of magnitude of raw bandwidth capacity, it must still be integrated into the IP-based infrastructure of the current and future Internet.

Reasons to Integrate

Beyond the need to accommodate the IP routing model, several technical and economic issues are driving efforts to integrate IP and optical technologies. For example, Internet service providers (ISPs) need the bandwidth multiplying capabilities afforded by solutions such as dense wave division multiplexing (DWDM) to meet anticipated network traffic demands. DWDM enables multiple OC-48 (2.5-Gbps) or OC-192 (10-Gbps) communications channels in the form of wavelengths or frequencies (termed lambdas and denoted with a "λ") to operate in parallel over a single fiber-optic cable.

ISPs might also find it less expensive to move large aggregates of IP packets in a purely optical format at transit points rather than the alternative: converting the optical signal to electronic format, processing at the IP layer, and converting back to optical for the next leg of the journey. Providers want to leverage IP traffic engineering mechanisms to rapidly configure optical-layer resources to address growing provisioning demands. Optical transport technologies also support strong protection and fast restoration techniques that ISPs must now consider as they carry more mission-critical data traffic.

Finally, some providers believe the functions offered by the asynchronous transfer mode (ATM) and synchronous optical network (Sonet) layers—long part of many network infrastructures—are no longer required if the same or better functions can be performed by a combination of the IP and optical layers. Eliminating these respective layers and their attendant hardware and operational costs is expected to lead to lower infrastructure costs and less complexity. Of course, this is not true in all cases—particularly for providers that continue to offer ATM and time-division-multiplexed (TDM) services—but for those that are banking on just handling IP traffic it might make sense.

Activities for better integrating IP and optical technologies are unfolding on many fronts. Router linecards supporting OC-192/STM-64 are in production, for example, and have been deployed in some networks. A new family of networking devices, loosely termed wavelength routers, are appearing from both traditional and start-up vendors. These devices run IP-like dynamic routing protocols and can attach and switch a large number of optical connections.

Within the IETF, numerous working groups are looking at better ways to run IP over and through optical networks. Most notably, the multiprotocol label-switching (MPLS) working group has proposed extensions to control optical cross-connects (OXC) and dubbed this multiprotocol lambda switching (MPλS).

Finally, other standards bodies and industry coalitions are working on standard interfaces to enable client entities (IP routers, for example) to signal for and establish optical connections through an optical transport network (OTN).[1] These groups include the Optical Internetworking Forum (OIF), the Optical Domain Service Interconnect (ODSI), and the International Telecommunication Union (ITU).

The IP-over-Optical Relationship

Before discussing recent architectural and protocol developments in the IP and optical space, it's useful to examine the functions performed at the IP layer and below. From there we can draw an evolutionary path toward the convergence of the IP and optical layers.[2]

Figure 1a illustrates what for many providers remains an abstraction: the functional layers operating in their networks. At the bottom is the optical layer, based on an OTN that consists of the following sublayers:

- The optical channel (Och) section defines an optical connection (light path) between two optical client entities. (The Och equates to a lambda in DWDM parlance.)
- The optical multiplex section (OMS) defines the connectivity and treatment for a multiplex or grouping of Och-level connections. (The OMS equates to a group of lambdas flowing over fiber-optic cable between two DWDM multiplexers.)
- The optical transmission section (OTS) defines how optical signals are transmitted over the optical media.

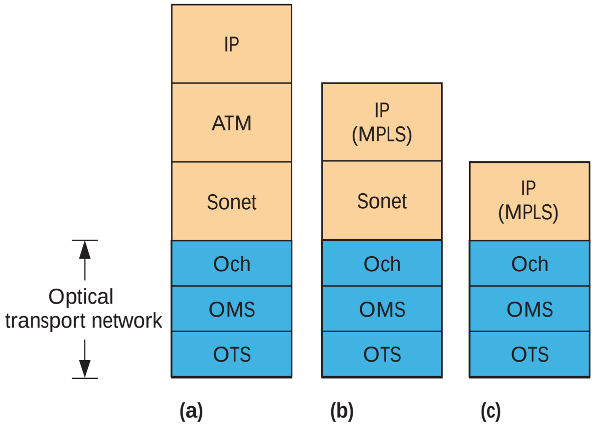

Figure 1. IP and the lower-layer evolution: (a) IP over ATM/Sonet/Optical, (b) IP over Sonet/Optical, and (c) IP over Optical.

Sonet layer. The Sonet layer takes lower-speed, TDM circuits (DS1 or DS3, for instance) from client devices (such as ATM switches), and places them in a synchronous frame format for transport across a higher-speed transport network (OC-3, perhaps). This function is typically performed by Sonet add/drop multiplexers (ADM)—so named because they can originate (add) and terminate (drop) tributary connections of varying speeds. ADMs are generally configured in fiber-ring topologies, and a Sonet network is formed by concatenating two or more rings using devices called digital cross-connects (DCS). Configuring an end-to-end TDM circuit can be a lengthy process because the provider must configure each ring and DCS along the path.

As part of its heritage of carrying multiple TDM (read, "voice") circuits in voice-carrier networks, Sonet supports a comprehensive suite of operations, administration, maintenance, and provisioning (OAM&P) functions, which are used to configure and manage circuits throughout the network. To protect against fiber cuts or other significant impairments, Sonet supports automatic protection switching (APS), which enables configuration of, and switchover to, a physically disjoint protect (backup) path in case of a failure along the working (primary) path. Service is restored very quickly (within approximately 50 milliseconds), but at the added expense of allocating bandwidth and equipment for the protect path.

ATM layer. The ATM layer, positioned above the Sonet layer, provides several powerful networking functions. It is a connection-oriented technology that requires the setup of a virtual connection (VC) between the source and destination before information can be exchanged. A VC can be built manually or automatically through a dynamic process that employs an ATM-specific signaling and routing protocol. ATM supports multiple service classes, which lets providers multiplex and transport data, voice, and video traffic with predictable performance. By defining a VC that traverses a specific path between two points in the network, providers can also use ATM to perform traffic engineering.

Like any networking technology, however, ATM has its trade-offs. Bandwidth efficiency is lost because ATM segments packets into 53-byte cells for transport—the price of the five bytes of control information in each ATM cell. Another issue is scalability: IP routing protocols do not scale well with numerous links. A VC is considered a link, and connecting N number of IP routers in a full mesh topology requires $O(N^2)$ VCs to set up and manage. Finally, ATM requires its own addressing scheme, routing protocols, and network management systems, which all add to network complexity and operational costs.

IP layer. On top of the ATM layer sits the IP layer, which encapsulates data, a fast-growing amount of voice, and some video into an IP packet, then forwards it hop by hop through the network. The IP layer provides an any-to-any, connectionless internetworking function and is self-healing in that packets can be dynamically rerouted around failed links, nodes, or networks.

Network Layer Evolution

Each of these functional layers performs the job it was designed for, but each was developed independently and with different technologies to serve different provider needs. Sonet, for example, was developed as a reliable transport for voice circuits, whereas ATM was created to support multiservice networks. It is not clear if these two layers will, or should, have a role in every network once all data is carried inside IP packets over big optical pipes.

There are inherent inefficiencies as IP packets move between the different protocol layers. ATM exacts the 5-byte cell tax, and due to its fixed bandwidth circuit orientation, Sonet may be under- or oversubscribed depending on the burstiness of the IP traffic flows.

Managing and operating the IP, ATM, and Sonet networks adds additional costs because separate network management disciplines, tools, and perhaps even staff are needed for each. Reporting and correlating different alarms and statistics among the networks introduces additional complexity. Moreover, the multiple layers may perform duplicate functions; for example, each supports some sort of recovery scheme.

IP over Sonet. Figure 1b shows the next phase in the evolution toward an optical Internet—one that is already supported in many networks today. In this depiction, ATM has been squeezed out in favor of running IP directly over Sonet, which eliminates the functional, operational, and cost overhead of maintaining a separate ATM network. This is largely possible because router technology has surpassed ATM switching in performance and capacity and because the IP router is the native forwarding vehicle for the dominant transfer unit: the IP packet.

In addition, adding MPLS to the IP layer provides two powerful enhancements. First, it enables traffic engineering through the ability to establish VC-like explicit paths through an IP-router network. Second, MPLS separates the control plane from the forwarding plane, which lets IP control protocols manage the forwarding state in devices that are not, and will never be, capable of identifying IP packet boundaries (such as ATM switches).

IP over optical. Figure 1c shows the target architecture that many providers are aiming for: the Sonet layer has been eliminated in favor of running IP directly over the optical layer. Eliminating the ATM and Sonet layers means fewer network elements to manage. A combination of enhanced IP and optical layer recovery, OAM&P mechanisms, and distributed routing provides fast restoration, fault detection, and provisioning. A new, lightweight framing structure, such as the digital wrapper, might replace the functions provided by Sonet framing for Och connections.[3] The existing suite of MPLS traffic engineering (TE) protocols has been extended to operate over optical networks, and the IP layer, particularly the IP routers, can now interface directly with the optical layer.

Routing at the Optical Layer

A future integrated IP-optical network will most certainly feature gigabit and terabit IP routers outfitted with high-speed optical interfaces, as well as DWDM transmission equipment of varying sizes and configurations. The wavelength router is another component that will be present in some cases. (It goes by several aliases, including optical cross-connect, OXC, optical switch, optical router, and lambda switch router. To maintain consistency with most literature, I use OXC hereafter.) The dynamically configurable OXC can switch optical-level signals received on an inbound port across a switch fabric to a corresponding outbound port. To be clear, an OXC does not route or switch packets; it deals strictly at the optical layer where the unit of transfer is a fiber or lambda.

An OXC contains N number of input ports and M output ports, each capable of switching a fiber or lambda at a particular bit rate (b), so the device's overall forwarding capacity is $N \times M \times b$. Internal to the OXC is a switch fabric, which can be electrical or optical. The electrical switch fabric has a proven track record in routers and switches. It enables fault detection and isolation by examining management information contained in the header of each transport frame (for example, Sonet frame headers contain this information in overhead bytes). Moreover, the conversion from optical to electrical and back (O-E-O) performed by OXC transponders allows clean-up of optical signals impaired by attenuation or dispersion (both of which must be overcome when transmitting light over large distances). On the downside, the O-E-O conversion adds cost and power requirements to the device. In addition, upgrading to a higher speed or a different signaling format might require a hardware upgrade.

A pure optical switch fabric does not perform O-E-O conversion as the signal passes from input port to output port, which means presumably lower costs and reduced power requirements. It also provides bit-rate transparency; that is, the fabric will switch data independently of the signal format or bit rate. The same switch fabric could theoretically be employed to support first OC-48, then OC-192, and later OC-768 traffic. The disadvantage is that without any electrical awareness, fault isolation could be challenging. Pure optical switch fabrics and complementary innovations such as wavelength conversion and optical-layer fault isolation are still in development but likely to appear in commercially available products beginning next year.

A dynamic routing protocol operates on each OXC and is used in networks of arbitrary size and topology (just as in IP networks) for neighbor discovery, topology awareness, path calculation, and forwarding state installation. Whereas IP routers forward packets on a hop-by-hop basis, an OXC can originate, terminate, or transit an optical connection. Optical networks are connection oriented, and a signaling protocol is required to set up and manage the connection-forwarding state along each OXC in the optical connection path.

Rationale for OXC

So why the need for an OXC? Doesn't this introduce another layer of complexity that the IP layer could perform? To some degree it might. Yet while the OXC's success and popularity depends, in part, on the availability, cost, and performance of future terabit IP router platforms, it remains attractive because it can provide certain useful functions:

- *More efficient forwarding of traffic aggregates.* IP edge routers will aggregate and multiplex IP traffic streams onto a single lambda, enabling more efficient switching in the network core than more granular operations at the electronic IP layer. This efficiency contributes to network scalability and lower costs.
- *Optical mesh topologies.* OXCs per-

mit mesh networks, which require fewer resources for protection and restoration than traditional Sonet ring networks. This is because idle or undersubscribed network resources (links and nodes, for example) can provide backup for multiple end-to-end connections. Mesh topologies also enable constraint-based paths between any two points in the network.

- *Optical bypass.* Transit traffic arriving at a provider's point of presence (POP) can be switched at the optical layer rather than demultiplexed into packets and processed at the IP layer. Expensive electronic IP router resources can be allocated for serving customer traffic originating or terminating at the POP rather than traffic that is just passing through.
- *Optical-layer reconfigurability.* By distributing topological awareness and signaling intelligence on the OXC, we can efficiently configure the OTN resource pool to provision services or react to failures. As the OTN carries more IP traffic, TE practices and protocols can drive configuration and provisioning of inter-router optical connections.

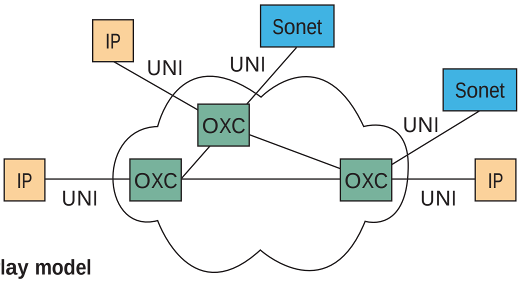

Overlay model

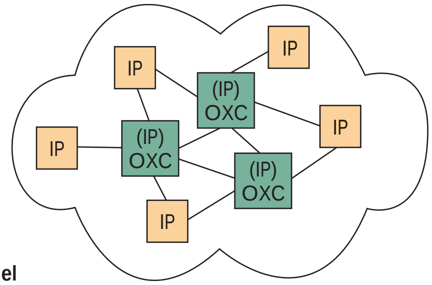

Peer model

Figure 2. Overlay and peer models. The overlay model supports separate routing protocols, addressing schemes, and network topologies between the client networks (such as IP and Sonet) and OTN. The peer model supports a single routing protocol, addressing scheme, and common topology across IP and optical devices.

Sycamore, Tellium, Cisco, and Calient are a few of the vendors working on an OXC-type device.

IP routers interconnected by static optical links could arguably perform these same functions, albeit with additional electronic processing. For transporting bulk IP traffic in optical format across a provider's backbone networks, however, it might make more sense to handle most core switching at the optical layer. IP routers at the network's edge could handle just the more complex aggregation and multiplexing functions.

Architectural Models

Several architectural models for IP and optical internetworking have been proposed in various forums. Not surprisingly, they are closely related to the models developed during the great IP-versus-ATM debates several years ago. One reason for this is that, like ATM, the OTN has been independent from the IP network; it has its own operational practices, network protocols, provisioning tools, network management applications, and so on. Also, like ATM, it is based on a connection-oriented or circuit paradigm as opposed to IP, which is connectionless. `Figure 2` illustrates the two basic models—overlay and peer—generating the most discussion.[4]

The overlay model has the following characteristics:

- IP routers and OTN OXC equipment are contained in two separate administrative domains.
- IP routers are attached to the nearest OXC over a user-to-network interface (UNI). A UNI implies that one side of the connection is a client (IP router) and the other side is a network (OXC).
- IP routers are unaware of the OTN topology. IP routers form adjacencies with each other over OTN-provided optical connections and exchange topology information about the IP network.
- IP network and OTN each run their own set of signaling and routing protocols (control plane), maintaining separate topologies and exchanging little or no topology information.
- IP routers can request (signal for) the OTN to establish an optical connection with other IP routers.

Providers who wish to keep the OTN and IP networks (or any other client network, for that matter) separate will prefer the overlay model. This could be for reasons of administrative control and because the OTN is offering and billing for circuit-based services connecting a variety of different client types (including IP routers, ATM switches, and Sonet ADMs).

The peer model has the following characteristics:

- IP routers and OTN OXC equipment are within a single administrative domain.
- IP routers and directly attached OXC neighbors form adjacencies to exchange topology information. Another (perhaps unfortunate) name for this peer-to-peer interface is the network-to-network interface (NNI).

- IP routers are fully aware of the OTN topology and vice versa. That is, all IP routers and OXC devices share a common view of the entire network topology.
- IP routers and OXCs run a common set of routing and signaling protocols and use a single addressing scheme.
- IP routers can request (signal for) an optical connection with other IP routers.

The arguments in favor of the peer model revolve around the fact that it is supporting a unified control plane for IP and optical network elements, and that it is optimized for IP-based services.

MPLS-over-Optical Network

MPLS is a logical choice for establishing a unified control plane, and it could be used to realize the peer model where IP routers and OXCs operate in a single administrative domain and maintain a single topology database. More specifically, we could extend the MPLS TE protocol suite to operate on IP routers and OXC devices.

Recall that the idea behind MPLS TE is to set up label switch paths (LSPs) across a network of label-switched routers (LSR) based on bandwidth or other policy constraints. MPLS TE components include a signaling protocol for setting up the LSP, a routing protocol (OSPF or IS-IS) with appropriate extensions for advertising the network topology and available link resources, and a mechanism for forwarding packets independently of their IP header and payload contents.

There are several analogs between an MPLS TE network and an OTN running OXCs. For example, LSRs and OXCs use a similar forwarding paradigm: switching information units from an input port to an output port. The LSR bases its switching on a label contained in each packet, and the OXC bases it on port number or lambda. Another similarity is that LSPs and optical LSPs are unidirectional point-to-point connections established through a path of (LSR or OXC) nodes that meets certain constraints. These similarities suggest that MPLS would be a sound choice for creating an open, interoperable control plane for integrated IP-optical networks.[5] The sidebar, "MPLS Traffic Engineering over Optical," examines the MPLambdaS (MPλS) efforts now under way in the IETF to extend MPLS TE over optical.

The number of distinct LSP-style connections traversing an MPLS-OXC network can be limited by label space, which is related in this case to how many lambdas can be packed onto a fiber. Current DWDM technology places this number at a couple hundred. Even with multiple fibers and wavelength conversion, that is orders of magnitude less than the 2^{20} labels available in IP routers (each packet's 4-byte label has a 20-bit label field). It might therefore be useful to aggregate multiple LSPs into one larger optical LSP to protect this limited resource and maximize the traffic that can leverage the relatively few very high capacity optical LSPs.

We can accomplish this first by using MPλS mechanisms to build an optical LSP between an ingress and egress IP router. This optical LSP forms a forwarding adjacency (FA) link, which the routing protocol's flooding mechanisms advertise into the topology database of all optical and nonoptical IP routers in the network. Any IP router in the network (even if not directly attached to the MPLS-OXC network) can consider the FA link in its path calculation when initiating an LSP-setup request.

When the LSP traverses the FA link, the IP router at the FA's head end employs label-stacking procedures to nest the smaller LSP inside the larger FA link for transit across the MPLS-OXC network. In this context, label-stacking means that the head end (ingress) router of the FA can direct labeled packets from many smaller LSPs through a single, larger optical LSP which is the FA link. In addition to making maximum use of available LSP resources, providers can apply policies to permit or deny certain routers from using these connections.

Optical UNI

The Optical Internetworking Forum is developing another approach, optical UNI (O-UNI), which provides an interface between optical clients and an optical network.[6] It defines a link (or links) for communicating control, signaling, and data packets between the optical client (IP router, for example) and adjacent OXC. The O-UNI was created so carriers could offer a simple, open, external interface for accessing circuit-based services across an OTN. The O-UNI defines only the interface and protocol interactions between the client and adjacent optical device.

An O-UNI can be *static*, which means that only control information (checking the identity and state of UNI links and devices) is exchanged between the client and OXC device. With a *dynamic* UNI, on the other hand, a source client generates signaling information, which it sends over the UNI to establish an optical connection with a target client. A client can be an IP router, ATM switch, Sonet ADM, or any optical client device that requests services from the OTN. The actions that can be invoked over the O-UNI include light-path creation, deletion, modification, and status inquiry. A client also might wish to register or deregister its identity with the adjacent OXC.

The O-UNI implies use of the overlay model in which the OTN topology is hidden from an overlay network of attached client devices. It may also be possible for a third-party entity (Network Management Station or NMS) to trigger the establishment of an optical connection on behalf of a source and target clients. The use of an NMS to control the provisioning of optical connections may be useful for providers who are familiar with a centralized approach to network management and provisioning.

While MPλS and O-UNI might seem like two distinct solutions, they are similar in that:

- both will employ the same protocol machinery (with extensions) being defined within the IETF, including RSVP and CR-LDP for signaling and the link management protocol (LMP) for link management and discovery between adjacent nodes, and

- both let the (source) edge device request (signal for) a dynamically established optical connection with a target device.

The basic difference between the two is the amount of information exchanged between client and OTN. The (implied) peer model of MPλS calls for routing protocol exchanges (topology information) between adja-

MPLS Traffic Engineering over Optical

IETF efforts now under way to extend multiprotocol label switching (MPLS) traffic engineering (TE) over optical networks are known as MPLambdaS (MPλS). Figure A, which presents a high-level view of an MPλS network, illustrates several important MPLS TE architectural and functional extensions being developed to support a combined network of routers and OXCs.

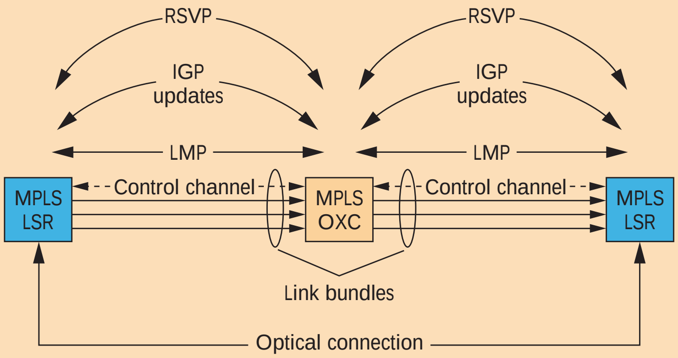

Figure A. MPλS network. An MPλS network consists of LSR and OXC devices interconnected by optical links. Protocols for advertising topology (IGP), verifying link connectivity (LMP), and signaling for connection setup (RSVP) operate over the control channel and enable the establishment of optical connections.

Link Bundles and Control Channels

For scalability reasons, a group of one or more unidirectional bearer-capable channels ("component links" in the form of fibers or lambdas) that interconnect a pair of MPLS-controlled OXCs (MPLS-OXC) (or LSRs) and an associated bidirectional control channel are collectively called and advertised as a single logical link. The control channel conveys only control information between adjacent MPLS-OXCs and might operate over a separate fiber, lambda, or even an out-of-band Ethernet connection. Other possibilities for establishing the control channel include embedding the control information in existing Sonet overhead bytes or using some form of subcarrier modulation (SCM).[1]

Link Management Protocol

LMP is a new control protocol that operates between adjacent MPLS-OXCs. It monitors control channel availability, verifies connectivity and availability of the component links, and supports fault isolation.

Routing Protocol Extensions

The routing protocol (OSPF or IS-IS) must be extended to encode and advertise characteristics of the optical connections. This information is used during the path calculation process to determine what links along a candidate path will satisfy the imposed constraints. The routing protocol must propagate information such as

- protection capability (if any) that the link configuration offers;
- encoding and bit rate of the link;
- whether the link is part of a group of links that will all be impacted if one is severed;
- any optical impairments, such as attenuation or dispersion, on the link that affect the optical signal's quality; and
- demultiplexing capacity of the link's receiving interface.

The last item determines what kind of optical connection a node's particular interface can terminate. For example, an edge router will advertise its optical interfaces as Packet-Switch-Capable, a Sonet ADM might advertise its interfaces as TDM-switch-capable, and a transit-only MPLS-OXC device might advertise its interfaces as fiber- or lambda-switch-capable. Optical connections can be established only between entities with similar link-multiplexing capabilities.

Signaling Extensions

The signaling protocols, resource reservation setup protocol (RSVP), and constraint-based routing using label distribution protocol (CR-LDP) convey label requests and label objects along an explicit path. The label's semantics must be extended to support not only packets but fibers, wavelengths, and TDM circuits. In addition, a link identifier is needed to indicate the specific component link in a bundle on which a label is allocated. Other extensions must let the signaling protocol set up bidirectional optical connections and request the use of one wavelength frequency end-to-end if no wavelength conversion is possible. Wavelength conversion enables an input lambda of a certain wavelength to be converted to an output lambda of a different wavelength.

Reference

1. D.W. Griffith and D.H. Su, "Standards: Standards Activities for MPLS over WDM Networks," *Optical Networks*, vol. 1, no. 3, July 2000, pp. 66-69.

cent devices and the O-UNI does not. The signaling information generated (and transmitted) by the source in the peer model is based on a complete view of the topology while the O-UNI requires only a topology subset. Furthermore, the O-UNI does not define control protocol exchanges between adjacent optical switches inside the OTN.

In general, IP-based service providers who own and manage their own router, fiber, and OTN infrastructure are more inclined to go with a peer model implementation. Providers who offer multiple services over their OTN will probably choose an overlay solution. Providers do, however, have the flexibility to use any combination. For example, a provider could configure a domain of IP routers and OXCs all running an MPLS or MPλS-based control plane for supporting ISP services.

Some OXCs might also support external O-UNIs that enable client devices to request and receive connection-oriented services. The routing and signaling of the O-UNI connection requests in the backbone would be handled by MPλS protocols. The provider could also use more traditional, centralized NMS-based tools for provisioning certain service-specific connections.[7]

Other Challenges

Much work and many challenges continue in this area. The way recovery can be coordinated across the IP and optical layers needs to be examined. Providers might not want to incur the costs of allocating extra OTN equipment and bandwidth for protect paths that might never be used. These tried-and-true mechanisms have set an industry benchmark of 50-ms restoration, however, and providers could be reluctant to fall back on IP layer recovery procedures that can take many seconds. MPLS fast restoration provides another option, but it adds some configuration overhead. Service providers might be content with restoration that takes several hundred milliseconds, on the other hand, because faster fault detection and improved routing algorithms could still achieve full restoration within the few milliseconds it takes a message to flow across the network.[8]

Another challenge is to lower the cost of optics. The OIF and vendors are currently investigating using very short-reach (VSR) optics to interconnect routers, OXCs, and DWDM gear over short distances. VSR optics are expected to be significantly cheaper than existing, commercially available short-reach optical interfaces. Several proposals are under discussion, but they basically all use vertical cavity surface-emitting lasers to drive one or more channels over a defined frequency to deliver an aggregate speed of OC-192 (10 Gbps). VSR optics can support Sonet and 10-GE (10x Gigabit Ethernet) interfaces, and many devices in provider POPs can benefit from this high-performance, low-cost interconnect technology. Commercially available VSR optics will be available in 2001.

Other efforts to evolve and adapt packet switching for the optical domain are under way, such as optical burst switching (OBS). This technique attempts to optimize the performance of packet bursts traversing an optical network.[9] Control or header information flowing over a separate control channel establishes a one-way path: no acknowledgment need be generated from the far end. A packet burst is accumulated and buffered in a fiber-optic delay line (actually a coil of fiber-optic cable) that provides a form of optical buffering similar to what random access memory does electronically. The burst is then sent over the reserved one-way optical path. OBS thus provides a degree of statistical multiplexing that utilizes the entire bandwidth for the duration of the burst.[10]

It is still early in the evolution of integrated IP and optical networking. But the industry is moving at "light speed" to develop standards and build network solutions that deliver on the promise of greater capacity, lower costs, and improved efficiencies. It will be interesting to check on the progress of the optical Internet in the not-too-distant future. ■

REFERENCES

1. *Recommendation ITU-T G.872, Architecture of Optical Transport Networks*, Int'l Telecommunication Union, Geneva, Feb. 1999.
2. D. Cavendish, "Evolution of Optical Transport Technologies: From Sonet/SDH to WDM," *IEEE Comm.*, vol. 38, no. 6, June 2000, pp. 164-172.
3. P. Bonenfant, "Optical Data Networking," *IEEE Comm.*, vol. 38, no. 3, Mar. 2000, pp. 63-70.
4. R.D. Doverspike et al., "Transport Architectures in an IP World," *Proc. IEEE Infocom 2000*, IEEE Press, Piscataway, N.J., 2000.
5. K. Bala et al., "IP over Optical Networks: Architectural Aspects," *IEEE Comm.*, vol. 38, no. 9, Sept. 2000, pp. 94-108.
6. N. Ghani, "Lambda-Labeling: A Framework for IP-over-WDM Using MPLS," *Optical Networks*, vol. 1, no. 2, Apr. 2000, pp. 45-57.
7. *OIF2000.125.1, User Network Interface (UNI) 1.0 Proposal*, Optical Internetworking Forum, July 2000.
8. O. Gerstel, "Optical Layer Signaling: How Much is Really Needed," *IEEE Comm.*, vol. 38, no. 10, Oct. 2000, pp. 154-160.
9. C. Alaettinoglu et al., "Towards Millisecond IGP Convergence," *Proc. NANOG Meeting*, Oct. 2000, Washington, D.C.; available online at http://www.nanog.org/.
10. C. Qiao and M. Yoo, "Choices, Features, and Issues in Optical Burst Switching," *Optical Networks*, vol. 1, no. 2, Apr. 2000, pp. 36-44.

Chris Metz is an IP architect for Cisco Systems. His current areas of interest include IP service differentiation, multicast, high-performance routing, and IP/ATM and IP/optical integration. He is coauthor of *ATM and Multiprotocol Networking* (McGraw-Hill, 1997) and author of *IP Switching: Protocols and Architectures* (McGraw-Hill, 1999). Metz is a member of ACM/SigComm and the IEEE. Prior to joining Cisco in 1998, he spent 14 years with IBM.