

Architecting the Services Optical Network

Eve L. Varma, Sivakumar Sankaranarayanan, George Newsome, Zhi-Wei Lin, and Harvey Epstein
Lucent Technologies

ABSTRACT

In this new millennium, the most valuable commodity will be information, as seen by the surge in demand for faster access technologies such as cable modems, xDSL, and wireless data. The tremendous growth in data traffic, particularly that associated with the Internet, is changing the way it is carried over public and private networks. Together with the rapid advances in optical networking technology and the spawning of a new category of wavelength services stimulated by new high-speed data requirements, this is dramatically changing network architectures and the relationship between network service providers and their customers. Incumbent carrier globalization and new carrier entry require support for a broader set of business models and range of service interfaces. The growth and expansion of networks imply a greater need for more scalable solutions and more automation for network maintenance. An intelligent optical core optimized for service optical networking, the *services optical network*, needs to take the best aspects of both the transport and data networking domains to meet the above needs. This article focuses on architecting the services optical network to meet the challenges of optical data networking, and includes a snapshot of supporting standardization activities.

INTRODUCTION

Transport networking has steadily grown more complex as a consequence of more sophisticated customer needs, the convergence of data and transport communications domains, the rate of traffic growth, and conditions imposed by external market and regulatory forces. Telecommunications increasingly relies on a diverse network of networks, with widely varying topologies, deployed services, and underlying business models. At the same time, traffic is growing at an unprecedented rate of 200 percent per year, while the capacity of silicon is only growing by

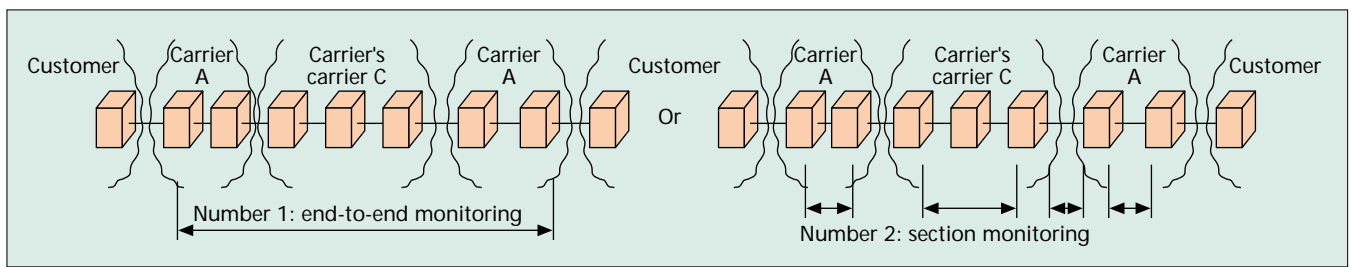
80 percent per year. Reliance on silicon alone cannot match this traffic growth.

Demand for access technologies such as cable modems, digital subscriber line (xDSL), and wireless data is rapidly increasing. Indeed, some analysts consider that data will eventually make up 99 percent of the content sent over telecommunications networks. Network operators have been finding that their more aggressive projections of traffic growth have tended to be wrong, and furthermore, that this demand for traffic occurs in often unexpected locations. The supporting network infrastructure needs to be inherently scalable to handle this traffic, and sufficiently flexible to respond to forecast uncertainties. At the same time, it is worth bearing in mind that the exponential growth in IP-based data traffic has not yet translated into a corresponding increase in revenue [1]. Higher-volume lower-revenue service mixes are driving the need for increased profitability reflected in near-term concerns with cost per bit, high-capacity fiber, and cost of spare equipment. It is clear that network operators will need to dramatically increase, and maximally share, backbone network infrastructure capacity. Moreover, lowering the cost per bit is not just equipment first cost, but requires efficient and optimized network design to reduce operation cost.

At the same time, the services offered, how quickly these services can be provided, and efficiency in supporting those services are key differentiators for network operators. Within the backbone/core network in particular, aggressive business models and limited network operator staff are driving the need for immediate availability, rapid deployment, and automated provisioning.

This growth of bandwidth and need to provide new sources of revenue to network/service providers have led to two major thrusts:

- Development of the infrastructure for the optical transport network (OTN), a scalable multiservice-capable and client-transparent transport infrastructure based on dense wavelength-division multiplex (DWDM)



■ Figure 1. A global multicarrier connection.

technology optimized for ultra-high growth, data-driven capacity demand (gigabit-level bandwidth granularity required to scale and manage multiterabit networks).

- Intelligent optical networking (service optical networking) — an optical control plane enabling a network that is self-managed with regard to topology/inventory discovery, route discovery, and connection setup. Such a network will support “on-demand” switched transport connections and service management that can support the range of operator business models.

This article focuses on architecting the services optical network to meet the challenges of intelligent optical networking, and includes a snapshot of supporting standardization activities.

OTN: THE FOUNDATION FOR THE SERVICES OPTICAL NETWORK

MOTIVATORS

As discussed in the introduction, incumbent carrier globalization and the entry of new carriers are creating market forces with a major impact on backbone infrastructure requirements. These factors also foster greater market segmentation and greater reliance on carrier cooperation. Service differentiation becomes an increasingly key factor, which requires verification of service level agreements (SLAs). Finally, the sheer size of networks implies a greater need for more scalable and automated network maintenance.

Hybrid solutions involving synchronous digital hierarchy/optical network (SDH/SONET) integration with DWDM technology have been increasingly deployed to tap into the full capacity of fiber plant, thus maximizing the return on existing facilities. This additionally allows service providers to flexibly add new services onto the existing fiber to accommodate new capacity demands. In these applications today, most of the transport networking functionality is provided by SONET/SDH systems that use DWDM systems. As network traffic grows and DWDM deployment continues, utilization of this approach for networked DWDM applications results in limitations in supporting the new multicarrier and multiservice networking requirements. Specifically, supporting *networked DWDM applications* using SONET/SDH layer functionality runs into three barriers:

- Limited support for SLA verification and fault sectionalization across multicarrier connections, even with operations systems involvement

- Inability to carry SONET/SDH clear channel (transparency)
- Non-scalable maintenance scenarios

SLA Verification and Fault Sectionalization

— Referring to Fig. 1, consider a global multicarrier connection between end customers involving an end-to-end global carrier network (e.g., with subnetworks within the United States and the United Kingdom) interconnected by a carrier’s carrier network.

Existing SONET/SDH network monitoring capabilities, called *tandem connection monitoring* (TCM), support one level of monitoring. This means that carrier A can monitor the end-to-end connection (case #1), or each carrier can monitor their individual parts separately (case #2). However, it is not possible to perform both types of monitoring. Thus, in addition to limiting what can be monitored, SLA verification requires tight intercarrier cooperation. For example, if the decision is made to provide end-to-end monitoring (#1), additional processes are required for fault localization. If the decision is made to provide section monitoring (any cascading of #2), an end-to-end carrier serving as the prime contractor of the service cannot determine overall signal quality. This carrier would therefore need to rely on customer complaints, or work to ensure tight coupling across carrier boundaries.


Transparent Carriage of SONET/SDH Services

— Referring to Fig. 1, suppose the customers are IP routers, and the service provided by carrier A is data transport between them. As an example, assume carrier A makes use of network facilities provided by carrier C, and desires to provide service protection via a SONET ring, as illustrated in Fig. 2.

For the SONET ring capability to be supported, nodes within carrier C’s network must not terminate the essential SONET overhead information (the entire OC-N payload including overhead must be transparently carried). This is not possible if carrier C employs SONET multiplexing, since carrier C would then terminate the SONET overhead information needed by carrier A. Thus, carrier A cannot build a ring across carrier C’s SONET network.

If carrier C uses a passive all-optical solution, carrier A has the needed SONET overhead, but then there is no overhead for carrier C to maintain its own network.

Maintenance Scalability — An interesting example of maintenance scalability issues is provided in the context of support for optically




■ Figure 2. Service protection via a SONET ring.


transparent subnetworks. Current SONET/SDH networks control faults by providing a specific alarm indication signal (AIS) indicating that the fault has been detected, and that downstream elements need not raise an alarm. Since it was assumed that SONET/SDH would always serve as the lowest network layer, no provision was made for an AIS between SONET regenerators. A new generic AIS signal has been defined in standards for use by DWDM systems to prevent downstream SONET/SDH network elements from alarming based on a DWDM line system failure. However, since a generic AIS can only be inserted at 3R points, there is no means of preventing alarm storms within all-optical subnetworks. This is illustrated in Fig. 3, which shows the effects of a cable cut causing many downstream nodes in various locations to generate loss of signal (LOS) alarms within an all-optical subnetwork. In a network with several cables per duct, dozens of fibers per cable, and hundreds of wavelengths per fiber, hundreds of thousands of LOS indications can occur, which will flood onto the management

communication network. Localizing the fault to the specific DWDM line system becomes very tedious, and requires that the network management system handle and correlate huge numbers of LOS alarms.

OTN Architecture — Recognizing that DWDM offered the means to support terabit per second line capacities, with associated gigabit per second path capacities, an aggressive work program was initiated to define the OTN within the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T). The OTN represents a new transport networking layer that is the next step beyond SDH/SONET in supporting data-driven needs for bandwidth and the emergence of new broadband services. It provides a multiservice-capable core infrastructure that leverages lessons learned from the SONET/SDH experience and adds optical technology to meet the challenges of the evolving telecommunications networking environment. This new layer is called the optical channel



■ Figure 3. The effects of a cable cut.



■ Figure 4. OTN connection monitoring in a multicarrier environment.

(wavelength), and it provides the gigabit-level bandwidth granularity required to scale and manage multiterabit networks and:

- Maximizes nodal switching capacity, which is the gating factor for reconfigurable network capacity
- Avoids very large numbers of fine-grained pipes that stress network planning, administration, survivability, management systems, and control protocols

The logical structure of the OTN networking interface, the optical transport module (OTM) [2], is described further below.

The optical channel (OCh) comprises an OCh data unit (ODU) and OCh transport unit (OTU). The ODU is a networkwide transport entity that can transparently transport a wide range of client signals including STM-16/64/256 signals, ATM streams, and IP/PPP and Ethernet signals after encapsulation with Generic Framing Protocol (GFP). It currently has three rates of approximately 2.5 Gb/s, 10 Gb/s, and 40 Gb/s, respectively, and is referred to as ODU k ($k = 1, 2, \text{ or } 3$).

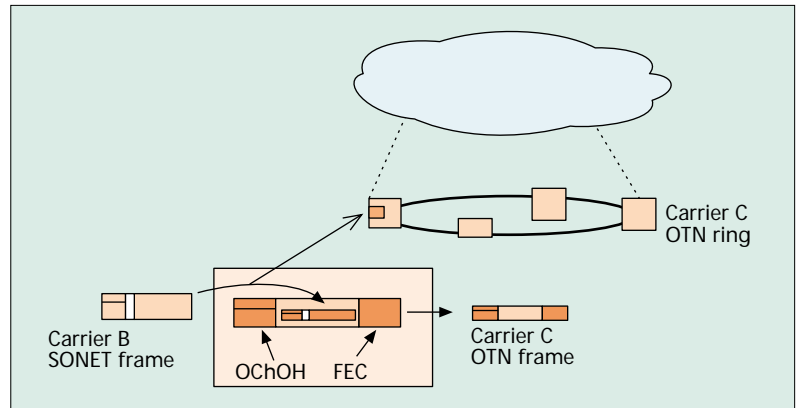
The ODU adds the overhead to support managed services in multi-operator DWDM-based optical networks in the client-independent manner essential for operating such networks. This overhead enables monitoring to support end customer, service provider, and network operator needs, providing for multiple levels of nested and overlapping connection monitoring. This overhead can also be applied to support protected domain monitoring, testing, and optical link connection monitoring. To condition the ODU for transport over a wavelength in a backbone network, it is transported within a frame that includes a forward error correction (FEC) code.

Overhead is also provided for single-channel and multiwavelength optical signals to support management of the “all-optical” parts of the OTN network. This overhead is typically transported via a separate optical supervisory channel wavelength.

OTN OVERCOMES THE BARRIERS

Now we revisit the DWDM networking barriers described in an earlier section, and see how they are resolved by OTN capabilities.

As illustrated in Fig. 4, OTN connection monitoring capabilities are optimized for a multicarrier environment, enabling fault localization and a wide range of SLA verification capabilities. Unlike



■ Figure 5. Supporting an OTN ring.

SONET/SDH, these capabilities can be deployed on a per-carrier basis (#4), end to end (#1), over several carriers (#2), at the same time.

The issues associated with supporting SONET ring capability in the multi-operator environment described previously evaporate when carrier C supports an OTN infrastructure. As illustrated in Fig. 5, when carrier C supports an OTN ring, the entire carrier A SONET frame is mapped into the ODU payload and transparently carried by carrier C, which uses ODU overhead in managing its network.

Maintenance scalability issues are also alleviated by the OTN OAM capabilities. Within the all-optical subnetwork illustrated in Fig. 3, OTN nonassociated overhead carried within the optical supervisory channel prevents alarm propagation. Thus, only the OLS adjacent to the fault will report a LOS alarm; forward defect indication (FDI) maintenance signals propagated to downstream nodes prevent LOS alarm flooding.

SERVICES OPTICAL NETWORKING

SERVICES OPTICAL NETWORKING OPPORTUNITIES

Services optical networking enables us to build on the optical transport infrastructure to provide switched connections and service management.¹ It enables an automatically switched optical network (ASON) that is self-managed with regard to topology/inventory discovery, route discovery, and connection management. This has major implications for carriers:

¹ Note that while this article focuses on OTN infrastructure and automatically switched optical networks, the architecture framework described is applicable to any switched transport network. In fact, the term automatically switched transport network (ASTN) was introduced to describe architecture and control plane aspects that are generic to any circuit-switched network.

Switched optical networking capabilities can provide an accurate, real-time view of the available bandwidth in the network as well as bandwidth utilization over time. This data can be used to fine tune network growth plans.

- Self-management will enable the introduction of new services and billing options for carriers.
- Greater automation of the service request process will reduce the delay between the time a service is requested, and when the connection is set up and billing can commence.
- The ability to support on-demand “Internet-time” connections will mean that carriers can provide short-lived connections with usage-based billing for broadband services. This enables provisioning of services such as bandwidth retailing by the hour, minute, or even second (bandwidth on demand) as well as bandwidth trading/reselling.
- By itself, automated route discovery can eliminate a manual step or handoff that can cost \$50–\$100 a connection.
- Route discovery combined with topology discovery can prevent connection route failures that occur frequently in current networks.
- Clients using a carrier’s network will be able to choose a class of service level appropriate for their needs (and pay just for what they need).

While the OTN will provide a transport infrastructure that supports self-healing networks, the switched optical network enables support for various types of survivability mechanisms, including fast mesh-based restoration of connections in addition to conventional options for unprotected and 1+1 protected connections. The shared mesh restoration option provides an opportunity to reduce bandwidth reserved for protection and provide more services on existing facilities.

Automatic topology, inventory, and route discovery provided by the switched optical network will enable networks to more fully utilize resources. Each facility will be discovered and may be used for connections as soon as it is deployed. Furthermore, since the network is the source of topology data, inventory cannot be lost due to administrative or reporting error as happens with current inventory systems. Anecdotal reports indicate errors in inventory databases relating to 15–30 percent of facilities being misconnected or missing. This results in a large percentage of connection routes failing to be usable. Switched optical network self-management allows the carrier to avoid these problems and allocate these resources to revenue-producing connections.

Finally, switched optical networking capabilities can provide an accurate, real-time view of the available bandwidth in the network as well as bandwidth utilization over time. This data can be used to fine-tune network growth plans, which are prone to be inaccurate in networks growing at 200 percent a year.

SERVICES OPTICAL NETWORKING BUSINESS ISSUES

It is essential that the switched optical network be architected in a manner that enables operators to bill for services they provide [3]. Thus, all reasonable business models must be supported. The business decisions associated with these models translate into network partitioning choic-

es, which in turn lead to explicit requirements for particular interfaces and information flows over those interfaces. Following are four examples of currently existing business models that can be used to facilitate the architecture and interface discussions of a later section [4, 5]. Note that one or more of these models might be used by various organizations of the same network operator.

The first business model to consider is that of an Internet service provider (ISP) that owns the network from the “ground up” (i.e., to the duct) and only delivers IP-based services. Here, there is no need to offer any other connection services on the infrastructure, and since it is fully owned there are no trust issues between the infrastructure provider and IP layers. This case is actually very unusual, since the rapid installation of capacity makes it unattractive for an ISP to actually own its entire infrastructure.

The second business model considered is an ISP that leases fiber or transport capacity from a third party, and only delivers IP-based services. As in the previous case, there is no need to offer any other connection services on that infrastructure. However, since the infrastructure is leased, there is a trust issue between the infrastructure provider and the ISP.

Now consider the business that provides this transport capacity to the ISP. In this third business model, the business owns the layer 1 infrastructure and sells services to customers who may themselves resell to others. There is a definite trust issue between this business and the client businesses.

Since the transport business cannot make any assumptions about the business of customers, it is highly likely that several customers may be engaged in the same business. If the customers are engaged in the IP business (as previously), there will be several instances of IP networks running over this network. However, customers may also engage in other data services as well (e.g., ATM); thus, a layer 1 infrastructure owner cannot make assumptions about what the client layer service may be without limiting their business opportunities. It may also play the role of a carrier’s carrier, where the customer network is likely to be a circuit network operated by another network operator.

The fourth business model that needs to be supported is that of a bandwidth broker. The subtle difference between this and the previous example is that the service provider may not own any transport infrastructure(s) that support those services, and the connection is actually carried over third-party networks. Here there are definite trust issues among the involved networks.

THE OPTICAL CONTROL PLANE ARCHITECTURE FRAMEWORK

The business models discussed above can now be used to derive the architecture requirements and set of interfaces to be supported by the switched optical control plane.

The most stringent architecture requirements are imposed by trust (or security) of the third and fourth business models in the preceding section. In order to realize the opportunities

described earlier, the switched optical network must have a means of discovering connectivity and capacity, a means of disseminating this information between peers, and a means of setting up a connection across the entire network.² If there are no trust boundaries in the network, the network user will be able to “see” the connectivity and capacity of each link in the entire network. From an operator perspective, this would involve a significant loss of business secrecy [5]. Indeed, with this information the user could set up end-to-end connections at will by commanding the individual network switches, and network operators would find it very difficult to bill for the service. To protect their business secrets, avoid loss of control over the usage of their resources, and allow the creation of billable services, it is necessary to restrict the amount of visibility and control granted to an end user. Similarly, it is also necessary to restrict the visibility and control granted to a peer operator.

These trust issues lead to significantly different interfaces:

- Between user and provider
- Between adjacent switches within a single provider domain
- Between providers

This is also the case for the Internet, as well as the current public switched telephony network (PSTN). In the PSTN network, topology information is largely manually entered. In the Internet, topology information is kept within a routing (Open Shortest Path First, OSPF) domain that usually corresponds to a portion of a service provider network. A different protocol (Border Gateway Protocol, BGP) interconnects domains (or operators). OSPF does not distribute routing information to IP clients; the Address Resolution Protocol (ARP) isolates clients from many routing decisions.

But how different are these interfaces? To answer this question, consider several issues related to naming and addressing, since these issues will lead to an understanding of which objects need to be transferred across each interface. The terms *name* and *address* are frequently used interchangeably, but there are specific and subtle meanings implied by these terms. The fundamental difference between a name and an address is that names do not have semantics of any kind, while addresses allow the object to be “located” in some space. Observe that a name has a longer lifetime than an address, a property that allows names to be portable over some useful space.

Not allowing users access to the addresses used by the switched optical network has value, since it allows the operators to organize their networks as they see fit, without affecting their users. The interface between the user and the provider therefore needs to talk about user names, not network addresses. This interface is frequently called a *user-network interface* (UNI), a term that may cause confusion because it is often unclear whether it means the physical interface between user and provider equipment, or a logical interface with properties as discussed above.

In a distributed connection management scenario, the interface between adjacent switches within a single operator (or trust) domain allows switches to exchange network topology informa-

tion in terms of capacity between the switches, as well as exchange signaling messages that will control the switches during connection setup. These actions occur at separate times, and can best be described as different service interfaces. The overall interface term used is a network-network interface (NNI). The same caveat applies: this interface is considered purely logical, and to comprise several distinct services. Since routing decisions are made on the basis of connectivity knowledge among all the switches within the domain, consider this an *interior NNI* (I-NNI).

Trust issues mandate that detailed topology information cannot be exchanged across the interface between two operator domains. Instead, the information exchanged is the set of users reachable across this particular link. Since user names do not have routing semantics, and we must have a bounded solution to an (in principle) unbounded network, we take advantage of the routing properties of network addresses and exchange reachability information in terms of sets of aggregated reachable addresses. When the addresses are well formed, this reduces to exchanging the set of subnetworks that are reachable from a link, rather than the set of all users that are reachable. The observant reader will notice that this information is nothing more or less than a topology description of the connectivity between the domains, rather than that between the individual switches. This interface has been called an *exterior NNI* (E-NNI), to distinguish it from the I-NNI discussed above.

The aforementioned switched optical control plane interfaces are illustrated in Fig. 6.

THE STANDARDS ROADMAP

OTN INFRASTRUCTURE

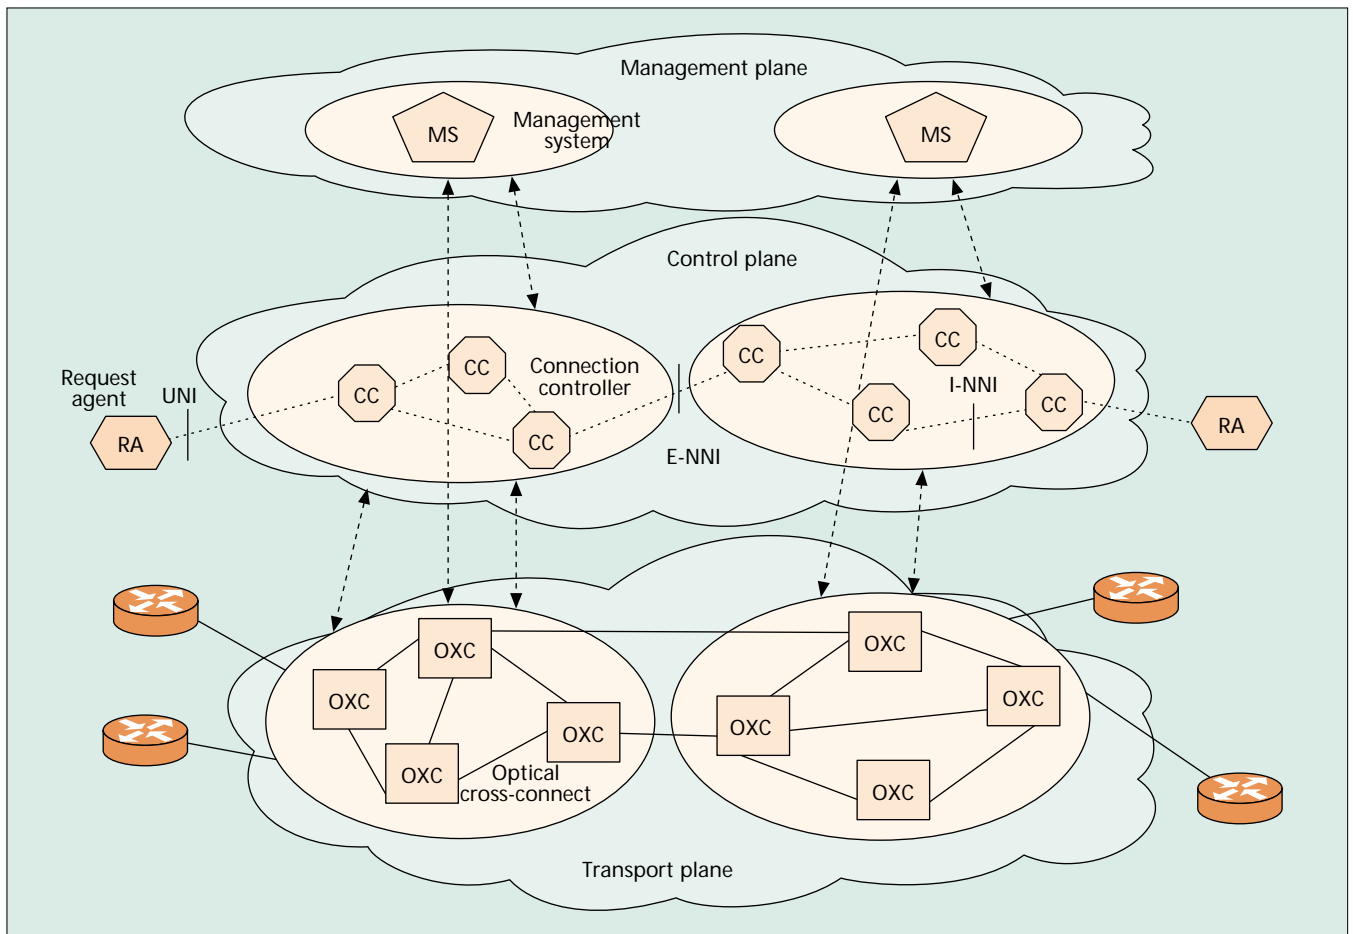
The ITU-T initiated work on OTN standardization by chartering a series of OTN Recommendations in 1997. This work commenced with G.872, “The Architecture of Optical Transport Networks,” which was stabilized during 1998, and formally approved February 1999. Work proceeded rapidly on G.709, “Interface for the optical transport network (OTN)” (mentioned earlier), which was stabilized during 2000 and approved at the February 2001 meeting of ITU-T Study Group 15. Equally important is G.959.1, “Optical Transport Network Physical Interfaces,” which defines optical characteristics, and was also approved in February 2001. Work is progressing on Recommendations dealing with management and equipment.

THE SWITCHED TRANSPORT CONTROL PLANE

Whereas OTN infrastructure standardization is primarily taking place within the ITU-T, this is not the case for control plane standards. Increasingly overlapping efforts among standardization bodies and industry fora are occurring because of the need for interdisciplinary expertise; for example, data networking protocol expertise of the Internet Engineering Task Force (IETF) and transport networking expertise of the ITU-T. To effect progress toward a coherent set of specifications that support operator needs, with requirements and architecture driving protocol solutions, a cooperative atmosphere is needed among the

Trust issues mandate that detailed topology information cannot be exchanged across the interface between two operator domains. Instead, the information that is exchanged is the set of users that are reachable across this particular link.

² The entire network is defined as comprising the individual networks of several operators working in concert.



■ Figure 6. Switched optical control plane interfaces.

various standardization bodies and industry fora. This is essential to avoid undermining the vision of a global switched optical network by inadvertently creating interoperability issues via conflicting requirements and specifications.

The current set of players involved in various aspects of switched transport networking specifications include the ITU-T, IETF, and Optical Internetworking Forum (OIF).

Within the ITU-T, work on the switched transport control plane as the ASTN umbrella of specifications (e.g., [6]) is focused within the following Study Groups:

- ITU-T Study Group 13, Question 10
 - Generic transport control plane architectural and functional requirements for global provider networks
- ITU-T Study Group 15, Questions 12 and 14
 - Optical control plane architecture requirements
 - Distributed connection management specifications
 - Neighbor and service discovery requirements, processes, and attributes

Within the IETF, two new Working Groups have been created to focus on the switched transport control plane, which has involved restructuring of related work efforts among other Working Groups and generalizing the multiprotocol label switching (MPLS) protocols as the generalized MPLS (GMPLS) umbrella of

specifications (e.g., [7]) for application to the control of circuit-switched networks.

- Common control and measurement plane (CCAMP)
 - Signaling protocols and measurement protocol specifications to support various technologies for connection control and discovery
 - Metrics, parameters, and properties that can be carried in protocols
 - Relationship between routing and signaling protocols
- IP over optical (IPO)
 - Switched optical control plane applications and requirements

Within the OIF, activities have been focused on establishing UNI signaling specifications based on modifications to Resource Reservation Protocol with Traffic Engineering (RSVP-TE) and Constraint-Routed Label Distribution Protocol (CR-LDP) protocols, for implementation as part of multivendor interoperability demonstrations.

CONCLUSIONS

The growth in demand for bandwidth and expectation of service fulfillment in “Internet time” require a different core network design for scalability and on-demand services. Within this article, we present an overview of a new broadband networking paradigm, services optical networking. The promise this paradigm offers is motivating industry momen-

tum toward establishing supporting specifications, reflected in rapidly accelerating OTN and optical control plane external standardization activities. Realizing the services optical networking promise requires unilateral commitment to a vision of multi-vendor and multi-operator interoperable networking that enables end-to-end switched services in a global environment. And a great deal of hard work!

ACKNOWLEDGMENTS

We wish to thank our colleagues Carmine Daloia and Maarten Vissers for their technical review and feedback, which has improved the quality of this article.

REFERENCES

- [1] A. Kalro, H. La Roche, and C. Roxlo, "Optical Networks for Data Applications: Market Drivers, Technology Trends, Network Architecture and Applications," ITS-2000, Sept. 2000.
- [2] ITU-T Rec. G.709, "Network Node Interface for the Optical Transport Network (OTN)," Feb. 2000.
- [3] OIF Carrier Group, "Carrier Optical Services Framework and Associated Requirements for UNI," T1X1.5 liaison doc., Oct. 2000, work in progress.
- [4] C. Brownmiller *et al.*, "Requirements on the ASON UNI," T1X1.5/2000-194, Oct. 2000.
- [5] "Considerations on the Development of an Optical Control Plane," Internet draft, draft-freeland-octrl-cons-01.txt, Nov. 2000.
- [6] ITU-T Rec. G.807, "Requirements for Automatic Switched Transport Networks (ASTN)," May 2001.
- [7] "Generalized MPLS — Signaling Functional Description," draft-ietf-mpls-generalized-signaling-04.txt, work in progress.

BIOGRAPHIES

EVE L. VARMA (evarma@lucent.com) is a technical manager with Lucent Technologies' Optical Networking Group. She received her M.A. degree in physics at the City University of New York. She has over 20 years of research experience within the telecommunications industry. Her primary research areas are currently focused on the automatically switched transport network (ASTN/GMPLS) and the optical transport network infrastructure (OTN). She is actively engaged in supporting the development of the associated specifications within global standards and industry fora spanning ITU-T, IETF, ANSI T1, and OIF. Between 1995 and 1997, she led the team responsible for designing and prototyping a distributed optical network management system as part of the Multiwavelength Optical Networking (MONET) Consortium (partially supported by DARPA). Her previous research experience includes characterization, analysis, and development of transmission jitter requirements; systems engineering and standards for SDH/SONET transport and network management applications; as well as associated enabling technology and methodology development/assessment.

SIVAKUMAR SANKARANARAYANAN (ssnarayanan@lucent.com) is currently a technical staff member with Lucent Technologies, part of Lucent's Optical Networking Group. He holds an M.S. degree in electrical engineering from Polytechnic University, Brooklyn, New York. His primary research areas are focused on the development of the architecture, and associated neighbor/service discovery and distributed connection management protocols, for ASTNs. His active research areas also include the network simulation of sig-

naling protocols in ASTNs, and the simulation of network performance under various fault restoration scenarios in MPLS networks. His previous research experience includes the development of protocols for synchronization distribution in SONET/SDH networks, systems modeling, robust and adaptive nonlinear control design, and the analysis of transport network architectures using formal specification techniques. He is actively contributing to the development of several key international standards in the area of ASTN, also serving as editor for the draft ITU-T Recommendations on auto-discovery, G.disc. In addition, he is a prolific contributor to ANSI T1, IETF, and OIF.

GEORGE NEWSOME (gnewsome@lucent.com) is a distinguished member of technical staff with Lucent Technologies, part of Lucent's Optical Networking Group. He received his B.Sc. degree in electrical engineering from University College, London, United Kingdom. He has over 30 years of multidisciplinary expertise encompassing software and system architecture analysis and definition, systems engineering and development of transport equipment and network software, CMOS customer ICs, application of microprocessor technology, technologies underlying design and composition of multi-usable software components, and methods of domain analysis. His primary research activities currently encompass optical networking architectures and automatically switched transport networking, including global standards development. Previous research activities include the establishment of formal specification techniques for transport, information modeling for network management based on open distributed processing, and systems and software engineering. He was a key software architect in the design and prototyping of a distributed optical network management system as part of the MONET Consortium. He is currently an active contributor to global standards on the ASTN and optical transport infrastructure. He has authored over 50 technical standards contributions during his standards tenure.

ZHI-WEI LIN (zwlin@lucent.com) is a member of Lucent Technologies technical staff, part of Lucent's Optical Networking Group. He holds an M.S.E.E. degree in electrical engineering from Cornell University, Ithaca, New York. He has multidisciplinary expertise on network survivability strategies, and network planning, architecture, and design. This encompasses a wide range of areas from ATM OAM work, to SDH MS-SPRING protection protocol design, to the G-MPLS/ASTN control plane specifications. From 1995 he worked at Telcordia as a senior consultant, where he was responsible for multitechnology network design and deployment, evolution planning, business plan development, and associated analysis and modeling. In addition to leading teams in these areas, he also played a lead role in representing Telcordia and the regional Bell operating companies within national and global standards. Since joining Lucent July 2000, he has expanded his standards activities to IETF and OIF as well. During the past 5 years, he has provided over 50 technical contributions to these standards and industry fora, many of which have been used as input toward approved international Recommendations. He is currently leading efforts in both ITU-T and IETF in specifying the G-MPLS/ASTN signaling protocol, serving as editor of draft ITU-T Rec. G.dcm (Distributed Connection Management).

HARVEY EPSTEIN (hepstein@lucent.com) is a technical staff member with Lucent Technologies, part of Lucent's Optical Networking Group in North Andover, Massachusetts. He holds a B.A. degree from Brooklyn College, New York, and M.A. and Ph.D. degrees in mathematics from the University of Wisconsin. He was responsible for operations architecture planning for optical networking, including optical UNI interface definition and contributed to development of ODSI Signaling Control Specification, which was the basis of the ODSI interworking demonstration in December 2000.

Realizing the Service Optical Networking promise requires unilateral commitment to a vision of multi-vendor and multi-operator interoperable networking that enables end-to-end switched services in a global environment. And a great deal of hard work!