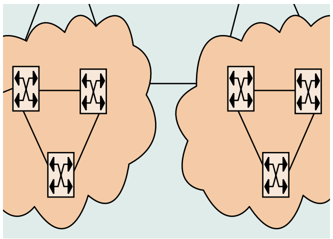

THE OPTICAL NETWORK CONTROL PLANE: STATE OF THE STANDARDS AND DEPLOYMENT

Debanjan Saha, IBM Research
Bala Rajagopalan, Tellium, Inc.
Greg Bernstein, Grotto Networking

ABSTRACT

In the last three years, the notion of an optical control plane has rapidly ascended from being a mere concept to a detailed set of protocol standards developed with broad industry participation. In this article we present a brief overview of optical control plane architecture and the associated protocols. We then examine the business drivers and inhibitors behind the optical control plane effort, the current state of the standards, interoperability status, and the open issues that need to be resolved before widescale deployment of this new technology can begin.

INTRODUCTION

In the last three years, the notion of an optical control plane has rapidly ascended from being a mere concept to a detailed set of protocol standards developed with broad industry participation. Some may say that the standardization effort has been too successful, since multiple prominent standards bodies have produced specifications with different assumptions on network operations. This raises several questions for the outside observer: What exactly is being standardized with regard to the optical control plane? Which standards bodies are involved? What are their objectives? Are there any commonalities in the various approaches taken? What are the differences? Most important, are there any issues in deployment of the optical control plane? This article attempts to answer these questions.

Specifically, we first describe the control plane architecture and the components being standardized. We then examine the business drivers and inhibitors behind the optical control plane effort, the current state of the standards, interoperability status, and the open issues that need to be resolved before widescale deployment of this new technology can begin.

CONTROL PLANE ARCHITECTURE

AN OVERVIEW

The notion of a "standard" control plane is to facilitate interoperability between multivendor equipment. But at which points in the optical network interoperable procedures are required must be determined. In this regard, the notion of a *control domain* is useful. A large optical network may be partitioned into smaller control domains mainly for two reasons:

- To enforce administrative, management, or protocol boundaries
- For greater scalability of the control and management planes

An example of the former is a multivendor optical network, which may consist of several interconnected vendor-specific control domains (or islands) with their own administrative and control procedures. An example of partitioning based on both reasons is the division of a global-scale optical network into a hierarchy of control domains based on geography and administrative boundaries.

From a control plane point of view, a partitioned network presents certain requirements. First, control procedures within different control domains may be heterogeneous. This could especially be the case in a multivendor network. A control plane spanning the entire network must accommodate the operation of heterogeneous control procedures within individual control domains. The ultimate goal, of course, is to be able to provision and maintain network connections across multiple control domains. Second, the scope of the overall control plane must be carefully prescribed. Finally, the implementation of a control plane requires information transfer between entities that participate in the control process. Partitioning invariably results in the curtailment of the information flow across control domain boundaries. This is also referred to as *information abstraction*. The precise mechanics of infor-

0163-6804/03/$17.00 © 2003 IEEE

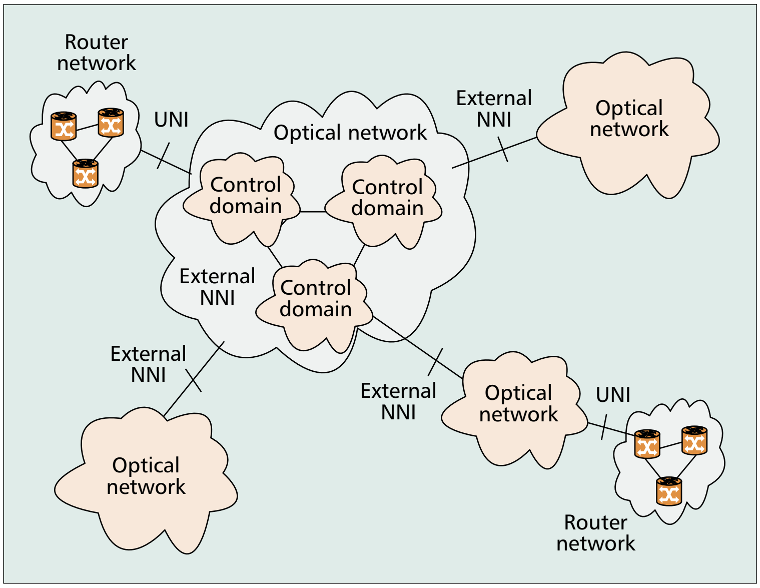

FIGURE 1. Optical control plane functional model.

mation abstraction determine the characteristics of the control plane procedures, and hence must be determined carefully.

Figure 1 illustrates the commonly accepted view of control interfaces in a network. Several optical networks, presumably administered separately, are shown. Each optical network is partitioned into control domains, and there are two client networks (router networks) connecting to the optical network. The control interfaces shown are:

- The user-network interface (UNI): This is the control interface between a node in the client network and a node in the optical network.
- The internal network-network interface (I-NNI): This is the control interface between two nodes within the same control domain.
- The external network-network (E-NNI): This is the control interface between two nodes in different control domains.

Internal node-to-node interfaces are not explicitly shown in the figure.

The control interfaces in Fig. 1 indicate the points where control plane interactions occur. An issue is whether the control plane functionality must be present in the network elements themselves. Typically, this is the case in IP networks, where the routing and signaling protocols are implemented directly in the routers. In the case of optical networks, however, the control functionality can be distinct from the transport functionality. Indeed, the control plane functionality may be implemented outside of the network elements by a proxy agent. Figure 2 illustrates the generic control plane abstraction, where the control agents depicted can be either built into or physically separate from the network elements. In the latter case, a single control agent can represent multiple network elements. The (internal) interface between the network elements and control agents need not be standardized.

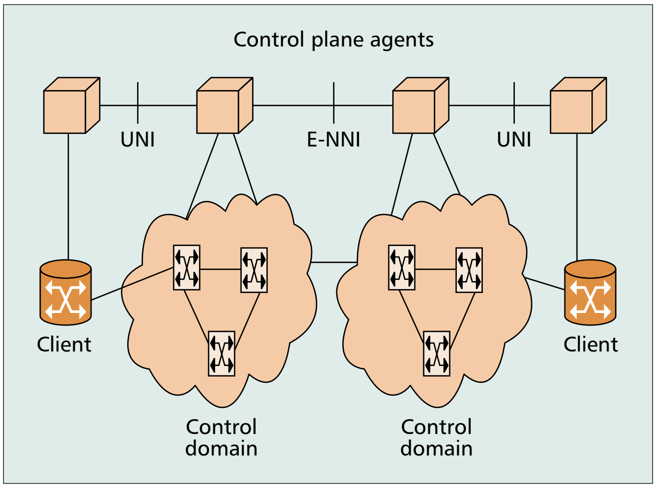

FIGURE 2. Control plane abstraction.

Whether the control functionality is integrated with the network elements or not, a function that can only be executed over the transport links between the network elements is neighbor discovery. This is described next. Also, the connectivity shown in Fig. 2 between the control agents defines the control plane *adjacency*. Two adjacent control plane agents in reality do not necessarily need direct physical connectivity. The only requirement is *reachability* between adjacent control agents. This reachability could be provided by any data communications network (DCN), including a direct physical link.

In addition to network partitioning, links may be aggregated or partitioned. Aggregation of simple links is called *link bundling*. A simple link in a bundle is also referred to as a *component* link. A bundled (compound) link may itself be a component link in another bundle. Thus, bundling could be recursive. The levels of bundling depend on how capacity information is represented in the control plane. Figure 3 illustrates different ways of bundling links. If the links are bundled into two link bundles, bundles 1 and 2, it is possible to retain the information that the two bundles go over two different conduits and potentially two different physical routes. This information is useful while computing route diverse paths for primary and backup paths for a service. If, on the other hand, for the purpose of most concise representation all the links are aggregated into a single bundle, bundle 3, the information about conduit diversity may be lost.

CONTROL PLANE FUNCTIONS

The main control plane functions are summarized below.

Neighbor discovery: In simple terms, neighbor discovery is a function whereby a network element automatically determines the details of its connectivity to all its data plane neighbors. These details include the identity of the neighbors, the identity of the link terminations, and so on. Neighbor discovery applies to both the UNI and the various NNIs shown in Fig. 1.

Routing: Routing broadly covers two control aspects. First is automatic topology and resource discovery; second is path computation. The first function enables the control agents to create a local view of the data plane connectivity and resource availability in the network. This procedure typically involves a mechanism to propagate the link connectivity information pertaining to a network element to all control agents in the network. Which information is propagated and how the information is represented in each control agent depends on the type of routing scheme. Path computation is a procedure whereby a control agent determines a path for a connection using the available topology and resource information.

Signaling: Signaling denotes the syntax and semantics of communication between control agents in establishing and maintaining connections. Signaling involves the use of standard communication protocols across the UNI and NNI.

Local resource management: This refers to the representation and accounting of locally available resources controlled by a control agent. Concise resource representation is essential for the scalability of routing mechanisms.

THE DATA COMMUNICATIONS NETWORK

The DCN refers to the communication infrastructure used for messaging between control agents in an optical network (or control domain). The DCN may also be used to provide connectivity between the control agents and element or network management systems. Typically, control communication is packet-oriented (e.g., using TCP/IP). Thus, the DCN is typically packet-based (e.g., IP) over an underlying network technology (e.g., asynchronous transfer mode, ATM). The DCN therefore provides network-layer (layer 3) connectivity between control entities in an optical network. Control communication may happen at several levels. For instance, communication is required between control entities in different control domains as well as control entities within the same control domain. The network infrastructure for interconnecting control agents across control domains could potentially be different than that used within a control domain. The term DCN, however, will be used to refer to the overall network infrastructure for control communication.

The main issues with regard to realizing a DCN are:

Network layer protocol used: Since the DCN provides network-layer connectivity, a specific network-layer protocol (and technology) must be selected. From a practical perspective, the choices today are IPv4, IPv6, or open systems interconnect (OSI) Connectionless Network Protocol (CLNP). A DCN may support more than one network layer with a built-in gateway functionality (e.g., CLNP and IP).

Network infrastructure: The main choice regarding the network infrastructure (on which the DCN is built) is whether or not the network is associated with the optical data plane. With a synchronous optical network/synchronous digital hierarchy (SONET/SDH) data plane, certain overhead bytes can be used to realize packet communication between control elements. In this case, the DCN network infrastructure is said to be realized *in fiber*. On the other hand, a separate network technology not associated with the optical data plane can be used to realize the DCN (e.g., an external IP network). In this case, the DCN is said to be realized *out of fiber*. As mentioned earlier, an overall DCN could consist of different parts that are interconnected, for instance, an in-fiber part in one control domain and an out-of-fiber part external to that control domain.

Application requirements: The DCN may support multiple applications, such as signaling, routing, and communication with management systems. These applications may pose certain requirements on the DCN, such as latency, packet loss, and security. In general, the DCN must support low latency and low packet loss. In addition, the DCN must be secure and not vulnerable to external interference.

DIFFERENT STANDARDS: COMPETING OR COMPLEMENTARY?

There are three main organizations developing optical standards for the control plane: the Internet Engineering Task Force (IETF), Optical Internetworking Forum (OIF), and International Telecommunication Union — Telecommunication Standardization Sector (ITU-T). There are three main sets of standards coming out of these bodies: generalized multiprotocol label switching (GMPLS), optical UNI (O-UNI) (work on NNI has begun), and the automatic switched optical network (ASON). Participants and observers question whether these different standard bodies and different sets of standards are competing or complementary. Different operating modes of these organizations further aggravate this tension. Here the activities of the various bodies are briefly reviewed along with the status of the standardization work in progress.

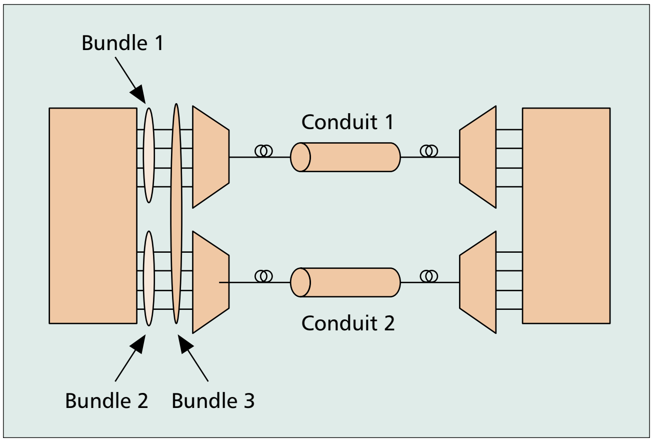

FIGURE 3 Link bundling alternatives.

THE INTERNET ENGINEERING TASK FORCE

The notion of an IP-centric control plane for optical networks was first described formally in an IETF Internet draft in November 1999. Note that this was after a number of vendors had already introduced the concept. This architecture was based on applying multiprotocol label switching (MPLS) control concepts to optical networks. It was first called multiprotocol lambda switching (MPλS), but later it was recognized that the same concepts could be generalized to control any circuit-switched network, including time-division multiplex (TDM), lambda-switch-, and fiber-switch-capable interfaces. Thus, the term generalized MPLS or GMPLS is now used to describe the application of MPLS protocols to controlling other networks. GMPLS-related work is carried out under the Common Control and Management Plane (CCAMP) working group.

Unlike MPLS, which has both a data plane and a control plane specification, GMPLS deals only with the control plane. The GMPLS architecture is described in an informational document [1]. Regarding optical network control, the GMPLS suite of protocols (Fig. 4) address the following aspects:

Link management: The Link Management Protocol (LMP) [2] is the IETF protocol specification for link connectivity verification, control channel management, link property correlation, and fault isolation.

Topology discovery: The IP link state routing protocols, Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS), have been extended with additional constructs to enable distributed topology discovery in optical networks [3, 4]. Work is ongoing on defining specific extensions for SONET/SDH networks [5].

Connection provisioning: The two MPLS signaling protocols, Resource Reservation Protocol with Traffic Engineering extensions (RSVP-TE) and Constraint Routed Label Distribution Protocol (CR-LDP), have been generalized for connection provisioning in optical networks. These are called GMPLS RSVP-TE [6] and GMPLS CR-LDP [6], respectively. Specific extensions to these protocols for supporting SONET/SDH networks have been specified [7]. Other extensions are expected as new requirements are brought up.

Connection protection and restoration: This is a relatively new area of work under the CCAMP working group. The aim of this work is to specify GMPLS RSVP-TE and CR-LDP based protection and restoration signaling mechanisms. Presently, restoration related terminology, analysis, and functional specification have been published [8, 9].

The GMPLS signaling specification and the RSVP-TE and CR-LDP extensions are currently being considered for elevation to proposed standards. The LMP and GMPLS OSPF and IS-IS extensions are in late stages of acceptance before being

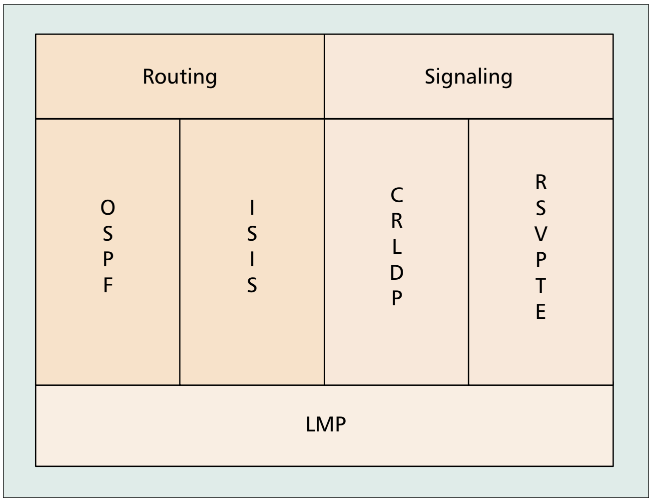

FIGURE 4. GMPLS protocol suite.

advanced as proposed standards. Once accepted as proposed standards, these specifications will be published as IETF Requests for Comment (RFCs).

THE OPTICAL INTERNETWORKING FORUM

The OIF is a consortium of optical networking vendors and service providers (carriers) whose goal is to expeditiously develop interoperability implementation agreements. Regarding the control plane, the OIF has mostly followed the functional models developed by the ITU-T in identifying the control interfaces. As per these models, the OIF work has focused on the interface between the user (or client) and the optical networks (the UNI), and the interface between optical control domains within a single carrier network (the NNI).

Optical UNI: The UNI is used by client devices to request optical network services, and the UNI implementation agreement was completed by the OIF in December 2001 [10]. This agreement is based on adapting the IETF GMPLS signaling specifications, notably GMPLS RSVP-TE, CR-LDP, and LMP (LMP used for neighbor and service discovery across the UNI), with certain customization for meeting specific carrier requirements. Work continues on version 2.0 of the O-UNI, with significant improvements in security, bandwidth modification, one-to-many call control, and several other features.

Optical NNI: The OIF NNI implementation agreement is currently being developed based on carrier requirements. These requirements stipulate a highly scalable and reliable control plane for international-scale optical networks. The NNI work is addressing the signaling and routing aspects, and the candidate solutions being considered are based on adapting the GMPLS signaling and routing protocol extensions with additional customization.

The OIF has done an excellent job in helping vendors and service providers reach an implementation agreement on UNI 1.0. This will facilitate deployment of at least a part of the optical control plane as more comprehensive standards such as GMPLS and G.ASON are worked on.

INTERNATIONAL TELECOMMUNICATION UNION

The ITU-T has been working on the architecture, functional models, and protocol specifics of the ASON. ASON is a client-server architecture with well-defined interfaces that allows clients to request services from the optical network (server). Within an ASON, signaling and routing interfaces between control domains implement the NNI functionality. A series of recommendations have been published by ITU-T pertaining to optical network control (Fig. 5). The salient ones are:

Network Architecture: G.8080: The architecture of ASON [11].

Neighbor Discovery: G.7714: Neighbor discovery [12].

Routing and topology discovery: G.7715: Architecture and requirements for routing in the ASON [13]. As with G.7713, it may be expected that routing protocol specifics will be specified in companion recommendations. It is very likely that IETF routing protocols will be adapted for this.

Signaling and connection management: G.7713: Specification of Distributed Connection Management (DCM) [14]. DCM refers to signaling for connection provisioning. G.7713 specifies the functional aspects. Specific protocol mechanisms aligned with the functional model are specified in G.7713.1 (based on the ATM private NNI, PNNI, signaling protocol), G.7713.2 (based on GMPLS RSVP-TE), and G.7713.3 (based on GMPLS CR-LDP).

Data Communication Network: G.7712: Architecture and specification of the DCN [15]. The DCN is the communication infrastructure used for control communications between optical network elements.

Due to the involvement of major telecom carriers, the ITU-T work is generally taken to be the baseline for functional requirements. The IETF protocol work is then adapted to meet these requirements, as seen from the work on DCM and routing above.

INTEROPERABILITY STATUS

Standards do not guarantee interoperability. Consider, for example, the BLSR/MSpring standard in SONET/SDH. Several years after being standardized, no two interoperable implementations of BLSR/MSpring from two different vendors can be found. The best way to ensure interoperability is to conduct trials involving multiple vendors to help uncover problems in early specifications and promote multivendor interworking. So far, trials have reflected varying degrees of maturity of the emerging control plane standards.

OIF UNI Interoperability Trial: In summer 2001 OIF sponsored a UNI interoperability trial involving 25 different vendors. The vendors underwent interoperability testing in the Interoperability Laboratory at the University of New Hampshire. Later on there was a public demonstration of interoperability at SUPERCOMM 2001 where the vendors successfully demonstrated UNI-N (network) and UNI-C (client) implementations used to dynamically provision service between clients and the optical network. It is worth noting that the reference model used for this interoperability event was a subset of what later became UNI 1.0.

This interoperability event was a valuable experience for the equipment vendors as well as the service providers. On the standards front it helped refine OIF UNI 1.0 specification. It also helped resolve some of the open issues in the IETF GMPLS signaling specifications. The OIF is considering conducting another interoperability event to test the emerging NNI specification.

IETF GMPLS Interoperability Survey: A large-scale multivendor interoperability trial for GMPLS interworking has not yet been conducted. However, a number of vendors have implemented various components of GMPLS and have conducted private interoperability testing. As a part of the standardization process, the IETF recently conducted a GMPLS implementation survey [16] where participants were asked to report the status of their implementation. A total of 25 participants including equipment vendors, software vendors, and service providers responded. From these responses it appears that GMPLS implementation is making steady progress,

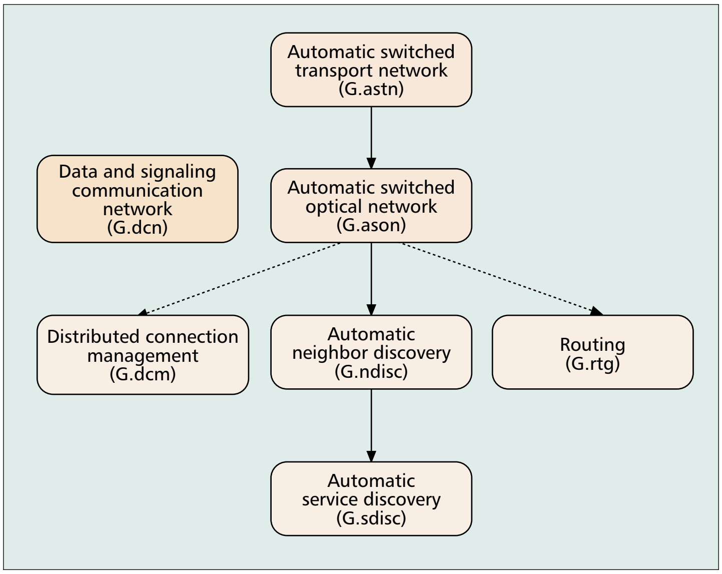

FIGURE 5. Different components of ITU-T defined optical control plane.

although the recent telecom downturn may have slowed the pace a bit.

Compared to the OIF UNI and IETF GMPLS standards, ITU-T ASON is still in its early stages. However, the protocol specifications under ASON are heavily leveraging existing standards and hence are expected to move quickly. Once defined, they can potentially reuse much of the code base from GMPLS and UNI, and be implemented quickly.

BUSINESS DIVERS AND INHIBITORS

In many ways, the optical control plane is a disruptive technology with the potential to revolutionize the transport infrastructure. In today's business environment, however, people have very little appetite for new technology unless it has a strong business driver. Hence, while there are differences of opinion regarding how standards should be defined and implemented, there is general agreement about why: it's all about making money! Service providers must increase their margins by reducing operational expenditures, and increase their revenue streams by introducing new services. The need to develop a common optical control plane to automate signaling and routing is clearly there. The introduction of agility into optical networks can reduce operational expenses and bring revenue-producing services online much faster. The top three drivers for a common optical control plane are:

New optical services: The capability to communicate between disparate network elements would enable new services and applications including bandwidth on demand, "time-of-day" bandwidth routing, optical virtual private networks (VPNs), new levels of QoS, the ability to add or drop services, and the ability to respond to changing needs quickly.

Operational cost savings: Development of a standard control plane can dramatically decrease operational overhead and cost. For example, standards-based protocols can reduce the testing and certification time for introducing new equipment into the network, hence saving time and reducing cost. It can also dramatically reduce time to market in provisioning even simple private line services. Benefits are quickly realized when private line provisioning intervals can be reduced from 45–60 business days to minutes, with human intervention only required at the ingress and egress nodes of the network.

While there are a number of business drivers, there are a number of inhibitors. It takes time and money to develop standards, and the number of talented resources required for the effort is diminishing worldwide. Loosely defined standards tend to be implemented in different ways by vendors; therefore, there is no "standard." This causes more problems than the standard was intended to solve. Inhibitors to the success of optical standards are:

Fragmentation of standards: A single standards body is not capable of addressing all aspects of an optical control plane. The IETF has taken the lead in defining the routing and signaling protocols. As GMPLS works its way into the ITU-T, there will be greater scrutiny, especially from the point of view of compatibility with existing transport network standards. Similarly, the ITU-T has little expertise in protocols that constitute GMPLS. Critical issues like billing, address management, restoration, and service-level agreements (SLAs) will have to be addressed before there is widespread adoption, and this level of expertise does not exist in a single standards forum.

Unclear business case: An optical control plane can potentially reduce cost and increase revenues, but the business case is still unclear. It will require a lot of investment to deploy this new infrastructure. Given the deep embedded base of legacy transport equipment, deployment will be gradual, and it is not clear how quickly service providers can recoup their investment. The worsening telecom market environment has further dampened enthusiasm for rapid deployment of new technology.

Clash of cultures. Any change faces challenges from those used to the status quo. The IP-cenric optical control plane technology is quite unfamiliar to the operators of transport networks, and its capabilities are unproven from their point of view. Also, it is not clear how this cross-technology standard will affect the organizational structure? As the network architecture collapses and becomes flatter, how will the various transport, switching, and IP organizations evolve?

The good news is that despite these inhibitors, the standardization process is moving forward. Also, many service providers have begun reorganizing the various operators and architects under a common management structure that is more receptive to cross-functional technologies, such as the optical control plane.

DEPLOYMENT ISSUES

The recent industry evolution has clearly altered the scope and timeframe of deployment of the optical control plane. In the early part of 2000, some predictions called for widespread deployment of intelligent optical networks containing hundreds or thousands of nodes, each with thousands or more of ports. However, the reality is that only recently a few networks with tens of nodes fitting the ASON model have started to be deployed. At the same time, some of the more sophisticated new applications that were expected to be widely deployed are slowly entering the trial and testing phase before final deployment. Several important issues need to be resolved before widescale deployment can happen.

Interoperability with legacy infrastructure: Today's transport infrastructure consists of hundreds of thousands of legacy SONET/SDH-based network elements. Any deployment of the optical control plane has to coexist with this legacy infrastructure and gradually evolve with it. It requires careful planning, extensive testing, and a well orchestrated evolution strategy to make sure that this transition is smooth and does not affect service.

Interoperability with management infrastructure: Current transport networks are supported by legacy operation support systems (OSSs) consisting of multiple layers or network management systems (NMSs) and element management system (EMSs). Any deployment of an optical control plane will require integration with the existing OSS and management systems.

Maturity of standards: Standardization of the control plane of intelligent optical networks is still far from over. Some components, such as provisioning, have received closer attention than others, such as restoration, and have reached higher levels of maturity. Some parts of the standards, such as multidomain and interdomain routing, are still in its infancy. All proposed standards need to receive more extensive testing and broad interoperability trials before final adoption and deployment.

Operational experience: Most of the protocols that constitute the optical control plane have their roots in IP networks. The operational environment and requirements of transport networks are very different from those of IP networks. Transport network operators will like to have more experience in using these protocols before they are confident enough to deploy them in the transport networks.

As a result of these conditions and current economic environment, we believe that the deployment of the optical control plane will be more evolutionary than revolutionary. We believe that operators will first deploy components of the optical control plane that give them the biggest "bang for the buck." We have already seen some examples of that in the more widespread acceptance of the O-UNI than the full GMPLS standards. We anticipate that even with GMPLS service providers will mostly likely deploy different components of it in an incremental fashion. One likely scenario is that link management (LMP) is deployed first, followed by signaling (RSVP-TE) for provisioning and restoration, and then routing and topology discovery (OSPF/ISIS).

CONCLUSIONS

In this article we reviewed the optical control plane architecture and the associated protocols being standardized by OIF, IETF, and ITU-T. We also discussed business drivers and inhibitors behind the optical control plane effort, interoperability status, and the open issues that need to be resolved before widescale deployment of this new technology can begin. We believe that deployment of a dynamic control plane in some shape or form is almost inevitable. Its long-term benefits far outweighs the short-term challenges. We expect that in the next several years standards will evolve to address some of the practical deployment issues. As the standards mature, the operators will get more exposure to this new technology and appreciate its benefits. Also, with time, data traffic will surpass voice traffic and the business case for a dynamic optical control plane will become stronger. So it is not a question of whether, but when you will see an optical control plane deployed in a network near you.

REFERENCES

[1] E. Mannie *et al.*, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," Internet draft, work in progress, April 2003.
[2] J. P. Lang *et al.*, "Link Management Protocol," Internet draft, work in progress, Mar. 2003.
[3] K. Kompella *et al.*, "OSPF Extensions in Support of Generalized MPLS," Internet draft, work in progress, Dec. 2002.
[4] K. Kompella *et al.*, "IS-IS Extensions in Support of Generalized MPLS," Internet draft, work in progress, Dec. 2002.
[5] E. Mannie and D. Papadimitriou, "GMPLS Extensions to OSPF and IS-IS for SONET/SDH Control," Internet draft, work in progress, Feb. 2002.
[6] L. Berger *et al.*, "Generalized MPLS — RSVP-TE Extensions," IETF RFC 3473, Jan. 2003.
[7] P. Ashwood-Smith *et al.*, "Generalized MPLS - CR-LDP Extensions," IETF RFC 3472, Jan. 2003.
[8] E. Mannie *et al.*, "Recovery (Protection and Restoration) Terminology for GMPLS," Internet draft, work in progress, May 2003.
[9] B. Rajagopalan *et al.*, "Generalized MPLS Recovery Functional Specification," Internet draft, work in progress, Jan. 2003.
[10] OIF, "User Network Interface (UNI) v1.0 Signaling Specification," Dec. 2001.
[11] ITU-T Rec. G.8080, "Architecture for the ASON."
[12] ITU-T Rec. G.7714, "Generalized Automatic Discovery Techniques."
[13] ITU-T Rec. G.7715, "Architecture and Requirements for Routing in the ASON."
[14] ITU-T Rec. G.7713, "Distributed Call and Connection Management."
[15] ITU-T Rec. G.7712, "Architecture and specification of the Data Communications Network."
[16] L. Berger *et al.*, "Generalized MPLS Signaling — Implementation Survey," Internet draft, work in progress, Oct. 2002.

ADDITIONAL READING

[1] B. Rajagopalan *et al.*, "IP over Optical Networks: Architectural Aspects," *IEEE Commun. Mag.*, Sept. 2000.
[2] E. Mannie *et al.*, "GMPLS Extensions for SONET and SDH Control," Internet Draft, Work in Progress, Feb. 2003.
[3] "Interim User Network Interface (UNI) Signaling Implementation Agreement for SUPERCOMM 2001," OIF contribution oif2001.152.4, May 2001.
[4] TeleManagement Forum, TMF 513, "Multi-Technology Network Management Business Agreement: NML-EML Interface v2.0"

BIOGRAPHIES

DEBANJAN SAHA (dsaha@us.ibm.com) is a well-known expert in the area of network protocols and systems. He is one of the principal authors of GMPLS standards in IETF and is a co-author of a forthcoming book on that subject titled *Optical Network Control: Architecture, Protocols, and Standards*. He has authored numerous papers and standards contributions. He is a co-recipient of IEEE Communications Society's William R. Bennett Prize Paper Award in the Field of Communications Circuits and Techniques. He is a frequent speaker at academic and industry events. He also serves as editor and technical program committee member of international journals and conferences. He is currently with the Network Software and Services Department at IBM T. J. Watson Research Center. Previously, he was with Lucent Bell Laboratories and Tellium Inc., an optical networking pioneer and world's first supplier of in-service intelligent optical networks. He has played a leadership role in a number of projects focusing on design and development of protocols for IP/MPLS routers, optical crossconnects, and Internet servers in both research and production environments. He received a B.Tech. degree in computer science and engineering from Indian Institute of Technology, India, in 1990; and M.S. and Ph.D. degrees in computer science from the University of Maryland at College Park in 1992 and 1995, respectively.

GREG BERNSTEIN (gregb@grottonetworking.com) is currently a consultant with Grotto Networking. Previously he was a senior technology director at CIENA Corporation looking over network control and management architectures after directing all software development at Lightera Networks (now Ciena's Core Switching Division) where his team applied signaling and routing techniques to the control of networks of Lightera switches (now the Ciena CoreDirector). Prior to the founding of Lightera Networks, he was a senior manager at Northern Telecom, specializing in new product development in the areas of ATM and IP. He has filed numerous U.S. patents in the areas of optical networking, packetized voice, congestion control, and encryption. He was previously involved in the development of a fiber-based digital video distribution system, then developed spread spectrum communication systems, and ran two DARPA sponsored advanced research contracts. He has published several papers appearing in *IEEE Transactions on Circuits and Systems*, *IEEE Transactions on Communications*, and IEEE INFOCOM and IEEE MILCOM conferences. He earned his Ph.D., M.Sc., and B.Sc. in electrical engineering and computer science from the University of California at Berkeley.

BALA RAJAGOPALAN (braja@tellium.com) is currently the principal architect at Tellium, Inc. He has worked extensively on IP-centric control of optical networks and optical internetworking architectures. He has also researched IP and wireless data networks for AT&T Bell Laboratories, Bellcore, and NEC, and he has made significant contributions to the IETF, ATM Forum, and OIF. He received his B. Tech degree from the Indian Institute of Technology, and M.S. and Ph.D. degrees from the University of Illinois, Urbana-Champaign, all in computer science.